

From: [FRUEH Terry](#)
To: [Henning, Alan](#); jeffrey.lockwood@noaa.gov; [SEEDS Joshua](#); [Kubo, Teresa](#); [Leinenbach, Peter](#)
Cc: [GROOM Jeremy](#)
Subject: FW: EPA/NOAA/DEQ review of ODF Riparian Rule write-up
Date: Wednesday, April 23, 2014 9:43:52 AM
Attachments: [Appendix 1.docx](#)
[Appendix 2.docx](#)
[Appendix 3.docx](#)
[Appendix 4.docx](#)
[Writeup2.docx](#)

Please ignore points 2 and 6 below – I made a mistake in copying text from another email without ensuring all of it was pertinent to you.

W. Terry Frueh
Monitoring Specialist
Oregon Dept. of Forestry
2600 State St., Building D
Salem, OR 97310
tel. 503-945-7392
TFrueh@ODF.State.OR.US

From: FRUEH Terry
Sent: Wednesday, April 23, 2014 9:30 AM
To: 'Henning, Alan'; 'Leinenbach, Peter'; 'Kubo, Teresa'; 'Jeffrey Lockwood - NOAA Federal'; 'SEEDS Joshua'
Subject: RE: EPA/NOAA/DEQ review of ODF Riparian Rule write-up
All,

Below is info from Jeremy concerning the model write up:

Attached is the methods section for the current analysis effort. I've tried to keep it from being too technical, so it is not written as tightly as a manuscript's methods would be. Also, it includes a fair amount of background material to assist with giving the current material some context. I wanted you to have this material a week prior to our meeting on the 21st. I thought it could aid us in our discussions.

The document is a beast (40 pages currently) with 4 appendices. The appendices have more technical material like exact model formulations, descriptions of FPA & state forest harvest simulation procedures, etc.

Things to keep in mind when reading it:

- 1) This is still firmly in a draft stage. I have not compiled all internal-ODF comments, and I'm still working on some edits of my own. The State Forest harvest procedure (Appendix 3) has been provided to state forests but we have not incorporated changes yet. I've made some changes to the FPA harvest scenario after receiving feedback from our Forest Practices experts (Keith Baldwin & Brad Knotts). Those changes have not percolated back to this draft yet either, but they are close.*
- 2) Next week my goal is to respond to earlier emailed questions from team members. I've been waiting to respond because I thought this document would help in answering questions.*
- 3) We are working on a literature review to give some context to the study's findings. So, this methods section just represents our own work on the analysis and does not yet provide context for the findings.*
- 4) We are planning on submitting this work for publication. Please be sure whoever receives this*

material to treat it accordingly.

5) Lisa Madsen has reviewed the document; she and I will also be meeting with two Weyerhaeuser statisticians (Jay Jones & Jack Giovanini) on the 23rd to review the analysis work to date.

6) We are still developing the agenda for our meeting on the 21st.

I look forward to meeting with you folks at 10 AM on May 5th at DEQ, room 6A.

Thanks,

Terry

W. Terry Frueh

Monitoring Specialist

Oregon Dept. of Forestry

2600 State St., Building D

Salem, OR 97310

tel. 503-945-7392

TFrueh@ODF.State.OR.US

From: FRUEH Terry

Sent: Thursday, March 20, 2014 12:51 PM

To: GROOM Jeremy; 'Henning, Alan'; 'Leinenbach, Peter'; Kubo, Teresa; 'Jeffrey Lockwood - NOAA Federal'; 'SEEDS Joshua'

Cc: ALLEN Marganne

Subject: FW: review of ODF Riparian Rule write-up

Folks,

I've got room 6A reserved at DEQ for our meeting on May 5 at 10 AM (note: us visitors will have to get a badge on the 10th floor to get back down to the 6th floor where the meeting is). I have the room until 1 PM in case we need it that late. Let me know if you have any questions, otherwise I look forward to seeing you then.

Terry

W. Terry Frueh

Monitoring Specialist

Oregon Dept. of Forestry

2600 State St., Building D

Salem, OR 97310

tel. 503-945-7392

TFrueh@ODF.State.OR.US

From: FRUEH Terry

Sent: Tuesday, March 18, 2014 9:10 AM

To: 'Henning.Alan@epamail.epa.gov'; 'Jeffrey Lockwood - NOAA Federal'

Cc: GROOM Jeremy

Subject: review of ODF Riparian Rule write-up

Folks,

Jeremy is working on the write-up of the model he is using to develop and test riparian rule prescriptions for ODF. Our intention is to get you this write-up the week of April 28th, and meet with you soon afterwards to address any initial comments and questions on the write up, while giving

you some time after this meeting to pull your input together. Please complete the attached doodle poll by this Friday to get this meeting on the calendar.

<http://doodle.com/a7qqb75rknverrrt>

Thanks,

Terry

PS – Alan, I had heard that you were taking over from Dave Powers about forestry issues in Oregon.

I'd be happy to bring you up to speed with this process if more info than what was presented at our January meeting would be helpful – just let me know. Also, if there are other people from EPA that should be included in this, please let me know.

W. Terry Frueh

Monitoring Specialist

Oregon Dept. of Forestry

2600 State St., Building D

Salem, OR 97310

tel. 503-945-7392

TFrueh@ODF.State.OR.US

Appendix 1 – Comparison of 40 day max with 7DAYMAX temperatures

In all RipStream temperature analyses, we are looking at the change in temperature between the upper temperature probe (2W) and the downstream probe (3W). This allows for annual differences in overall stream temperature. What we care about is the change (increase) in the difference.

The PCW examines a moving mean of the hottest daily temperatures taken across a seven day period (referred to here as 7DMM, or seven day moving maximum). The PCW considers all temperatures across the summer. If any one 7DMM is >0.3 C above what it should be due to human activity, the PCW is not in compliance.

We wanted to know the magnitude of the temperature change and related habitat variables. For this reason we developed the effects analysis. The temperature aspect of the effects analysis used a pre-defined 40-day period (July 15 – August 23). This standardized the timeframe of examination (unlike the PCW; see Table A1.1 for a list of analysis differences). To simplify the analysis, we used the mean of difference between the two probes' 40 daily maximum values (40-day max). So, for one year, one site, we obtain a single temperature value.

In comparing these two values, the first striking (yet obvious, in retrospect) feature is that the 40-day max is effectively the mean of the 7DMM values for those 40 days (Figures A1.1). Therefore, approximately half of all 7DMM values will be higher than the 40-Max and half lower. The 40-day max models (i.e., effects analysis temperature models) predict the mean response of 7DMM for a site but do not predict individual (e.g., maximum) 7DMM values.

A key point is that the PSTM and effects temperature models tell us what the predicted 7DMM are – on average – for a 40-day stretch in the summer. Is this value important? It appears that the PCW is concerned with the increase of any 7DMM temperature above 0.3 C. This does not mean that the warmest 7DMM are of concern any more than the coldest (e.g., Figure A1.2). What matters is if 7DMM are > 0.3 C above a no-impact variable. Modeling the mean response of 7DMM may be a reasonable approximation. That said, there can be a 2.5 C difference between the 40-day max and the maximum (Figure A1.3) or minimum (Figure A1.4) 7DMM value during the 40-day time period.

Table A1.1. Similarities and differences between temperature metrics and approaches of the PCW and Effects Analysis/PSTM models.

Difference	PCW	Effects analysis/PSTM models
Temperature measure	7DMM	40-Max (effects analysis includes mean, min, daily range)
# values per reach per summer	~60-90	1
Include site variables?	No	Yes
Utilize upstream control info?	Indirectly	Directly
Estimate magnitude?	Not advised	Yes
Evaluate PCW	Yes	Not with individual 7DMM

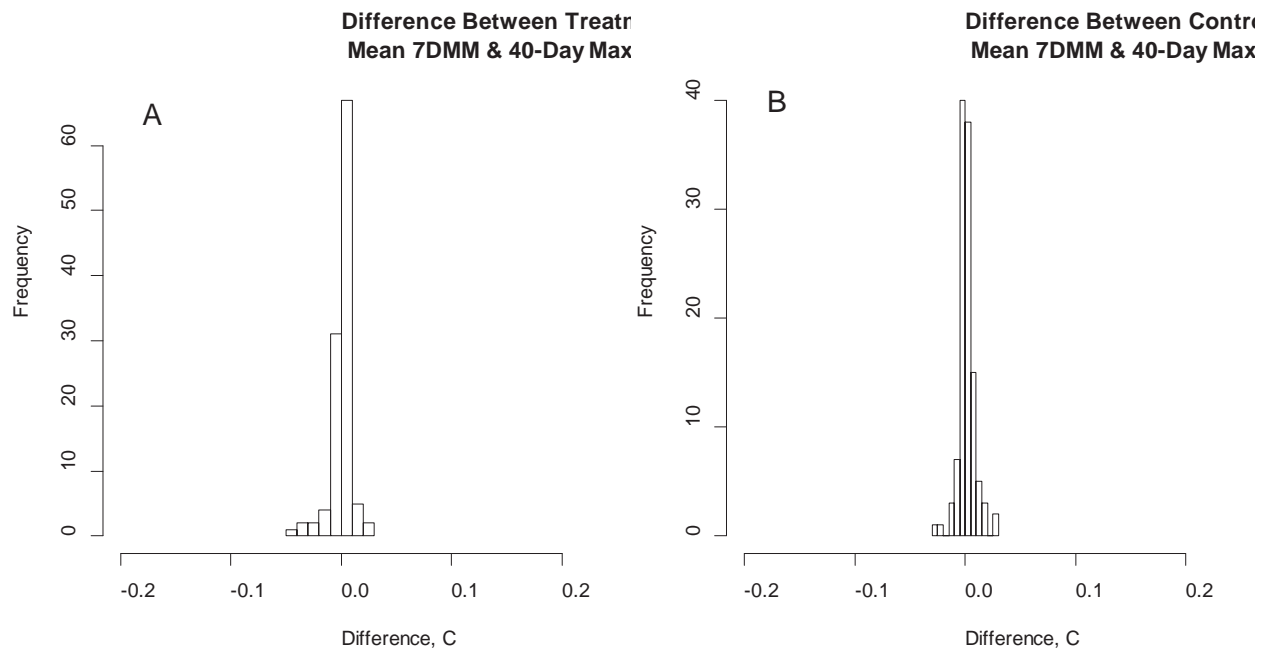


Figure A1.1. Histograms of the difference between mean 7DMM and 40-day max temperatures in the treatment (A) and control (B) reaches.

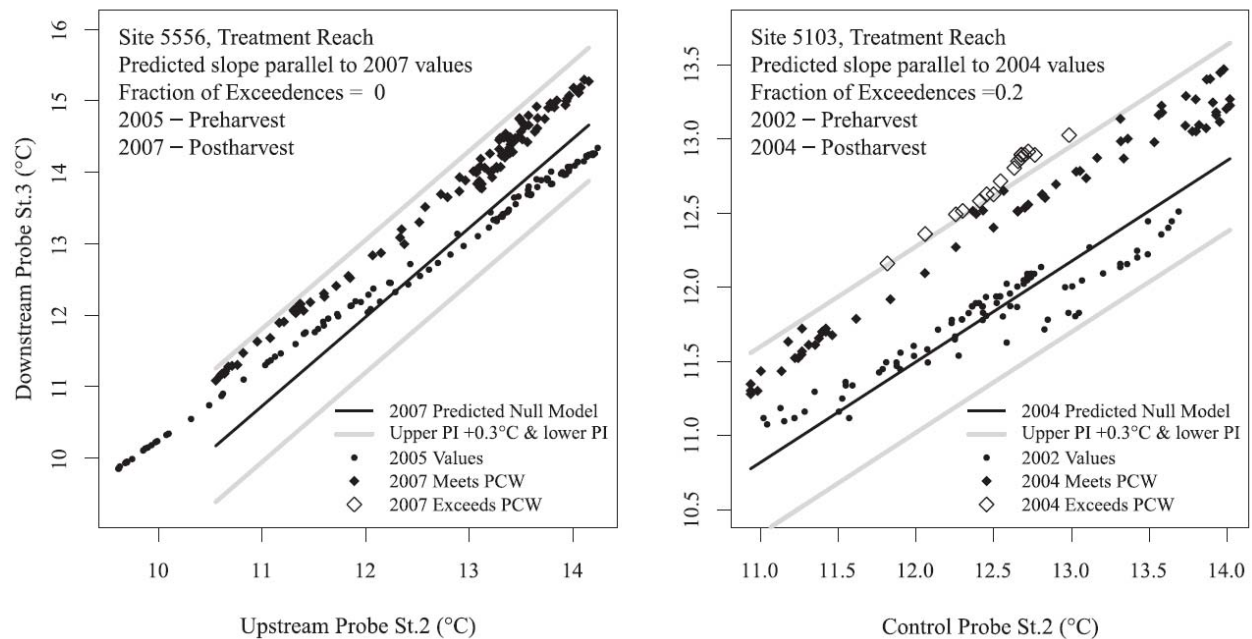


Figure A1.1 (From Groom et al. 2011a). Two examples of year-pair PCW evaluations. Each example is a preharvest to postharvest comparison of a single site's treatment reach using 7DAYMAX stream temperature data from 2W (upstream) and 3W (downstream). Black lines

represent null model predicted values for the postharvest data. The bottom and top gray lines represent the predicted null model's lower 95% prediction interval (PI) and an upper 95% PI +0.3 °C, respectively. Solid diamonds represent year 2 values that fell below the upper 95% PI +0.3 °C limit, and the larger open diamonds represent values above the limit (comparison for site 5103 represents a PCW exceedance).




Figure A1.3. Difference between 40-day max values and maximum 7DMM values for the same time period. Data are compared by site, year, and reach.




Figure A1.4. Difference between 40-day max values and minimum 7DMM values for the same time period. Data are compared by site, year, and reach.

Appendix 2 – Bayesian Parameterization of models

A2.1 Temperature Model

Below is the temperature model, constructed for a WinBUGS analysis. Code annotations appear after hashmarks(#). It is critical to note that Bayesian programs such as WinBUGS and JAGS do not process items in a linear order. The model and all components are run simultaneously. For this reason, an item defined at the top of the program is not estimated “first.”

The temperature model includes for a correlation between the slope (Control) and intercept random effects. Random effects are by site and are estimated as bivariate normal ($\mu = 0$).

Anything in the model that is followed by a [~] represents a parameter distribution for WinBUGS or JAGS to estimate (stochastic node). Anything followed by [<-] is a logical expression (logical node), likely containing manipulations of estimated parameters.

```
model {

## Setting Priors
  for (i in 1:ngroups)    #33 groups (sites)
  {
    alpha[i] <- B[i,1]    #Site intercept
    beta1[i] <- B[i,2]    #Site slope, by control reach temp
    B[i,1:2] ~ dmnorm(B.hat[i,], Tau.B[,])
      #multivariate (bivariate) normal distribution priors with per-site
      #values of i, precision Tau
    B.hat[i,1] <- mu.int    #fixed effect: intercept
    B.hat[i,2] <- mu.beta1  #fixed effect: for Control slope
  }

  mu.int ~ dnorm(0, 0.001)    # Hyperpriors for random intercepts
  mu.beta1 ~ dnorm(0, 0.001)  # Hyperpriors for random slopes (Control)

  Tau.B[1:2,1:2] <- inverse(Sigma.B[,])    #constructing the random effects
  Sigma.B[1,1] <- pow(sigma.int,2)          #variance-covariance matrix
  sigma.int ~ dunif(0, 100)                # SD of random intercepts, random uniform
  Sigma.B[2,2] <- pow(sigma.beta1,2)
  sigma.beta1 ~ dunif(0, 100)              # SD of random slopes, random uniform
  Sigma.B[1,2] <- rho*sigma.int*sigma.beta1
  Sigma.B[2,1] <- Sigma.B[1,2]
  rho ~ dunif(-1,1)                        #residual "random" effects
  covariance <- Sigma.B[1,2]

  beta2 ~ dnorm(0, 0.001)    #Fixed effect: TRLength
  beta3 ~ dnorm(0, 0.001)    #Fixed effect: Shade
  beta4 ~ dnorm(0, 0.001)    #Fixed effect: Grad1Q

  tau <- 1 / ( sigma * sigma)    # Residual
  sigma ~ dunif(0, 100)          # Residual standard deviation

# Likelihood
# n= number of temperature measurements (119)
```



```

for (i in 1:n) {
  response[i] ~ dnorm(mu[i], tau) # The 'residual' random variable, with
    #a mean of mu (below) and precision tau (above)
  mu[i] <- alpha[site[i]] + betal[site[i]]* c_ControlTemp[i] +
  beta2*c_TReachLength[i] + beta3*c_Shade[i] + beta4*c_GradlQ[i]
  #mu = equation 1. It is the expectation or predicted result given
  #parameter estimates and data values
}
}

# Initial values for model to operate with
inits <- function(){ list(mu.int = rnorm(1, 0, 1), sigma.int = rlnorm(1),
mu.betal = rnorm(1, 0, 1), sigma.betal = rlnorm(1), rho = runif(1, -1, 1),
sigma = rlnorm(1),
beta2=rnorm(1), beta3=rnorm(1), beta4=rnorm(1))}

# Parameters to provide estimates of
parameters <- c("alpha", "betal", "beta2", "beta3", "beta4", "mu.int",
"sigma.int", "mu.betal",
"sigma.betal", "rho", "covariance", "sigma")

# MCMC settings
ni <- 4000 #number of iterations
nb <- 2000 #number for burn-in (toss out first 2000 iterations)
nt <- 2 #Thinning: don't use every other value
nc <- 6 #Number of chains

# Start Gibbs sampler
out <- bugs(win.data, inits, parameters, "lme.GradShade.txt", n.thin=nt,
n.chains=nc, n.burnin=nb, n.iter=ni, debug = TRUE)

```

A2.2 Shade Model

The Bayesian formulation of the shade model is provided below. The model was run in JAGS. Code annotations appear after hashmarks(#).

The shade model is a weighted regression. Note that the values for `Logit_shade_var` provided for fixed sampling variance, used in constructing the weighted regression. The variables
`PctDiffBAlt100` = Percent Difference in Basal Area, less than 100' from stream
`Prel00PctHWD` = Percent hardwood within 100' of stream, pre-harvest
`TreeHt100m` = Tree height pre-harvest, within 100' of stream

Anything in the model that is followed by a `[~]` represents a parameter distribution for WinBUGS or JAGS to estimate (stochastic node). Anything followed by `[<-]` is a logical expression (logical node), likely containing manipulations of estimated parameters.

Within JAGS, we specified 20,000 iterations with 6 chains. By default, JAGS will discard the first 50% of iteration estimates and use enough to obtain 1000 estimates (1 in 10 thinning in this case).

```

#####SHADE#####
###priors for shade (suffix=S)
  alphaS ~ dnorm(0, 0.001)
  beta1S ~ dnorm(0, 0.001)
  beta2S ~ dnorm(0, 0.001)
  beta3S ~ dnorm(0, 0.001)
idelS ~ dunif(0, 100) # Residual
for (i in 1:n.sites)
{
  for (j in 1:2) #j = pre or post-harvest
    {tauSh[i,j]<-1/Logit_shade_var[i,j] #fixed variance (for weights)
    }
}

#### Likelihood
  for (i in 1:n.sites)
{
  muS[i,1]<-LogitShade[i,1] ##This is pre-harvest shade, using raw values

  Sh_post[i,2] ~ dnorm(LogitShade[i,2],idelS) # The 'residual' random
  ##variable, process variance (idelSh = "i-delta-shade")
  LogitShade[i,2] ~ dnorm(muS[i,2], tauSh[i,2] ) #setting up a fixed
  ##(weighted) sampling variance
  muS[i,2]<-alphaS + beta1S*PctDiffBAlt100[i] + beta2S*Pre100PctHWD[i] +
beta3S*TreeHt100m[i] ## muS = predicted shade value given model

}

##parameters for estimation
jags.params <- c("alphaS", "beta1S", "beta2S", "beta3S", "idelS")

##initial parameter values (randomized)
jags.inits<- function(){ list(alphaS = rnorm(1), beta1S = rnorm(1),
beta2S=rnorm(1), beta3S=rnorm(1), idelS = rlnorm(1))}

```

A2.3 Combined Model

The combined model consists of the temperature model from A2.1 and shade model from A2.2. However, there is a (logistic) transformation between the logit of shade values and the values for shade entered into the temperature sub-model here. Shade parameters have their own postscript, S (e.g., alphaS = intercept parameter for shade). Commentary on many parameters is not included below; see A2.1 and A2.2.

For predicting stream temperature changes, the (derived) variables `delta_sim`, `delta_muT_sim`, `estTemp[i]`, and `estTempNE[i]` are key. `Delta_sim` is the across-site mean of all predicted temperature change. `Delta_muT_sim` is a per-site estimate, calculated by subtracting the two harvest scenarios `estTemp` (harvest) and `estTempNE` (NE = no effect, or no harvest). The simulated value for Percent Difference in Basal Area is `PctDiffBASim`.

Following priors, the shade model is followed by prediction steps (relying on temperature parameter distribution estimates) and finally the temperature model. Running this model using

JAGS with 600,000 iterations appeared sufficient to reach a stable answer (300,000 iterations discarded as burn-in, 1 in 300 retained for estimates).

```

#### Priors ####

## Priors for temperature
for (i in 1:n.sites){
  alpha[i] <- B[i,1]                #Site intercept
  beta1[i] <- B[i,2]                #Site slope, by control reach temperature
  B[i,1:2] ~ dmnorm(B.hat[i,], Tau.B[,])
  B.hat[i,1] <- mu.int
  B.hat[i,2] <- mu.betal
}

mu.int ~ dnorm(0, 0.001)            # Hyperpriors for random intercepts
mu.betal ~ dnorm(0, 0.001)         # Hyperpriors for random slopes (control
                                   # reach temp)

Tau.B[1:2,1:2] <- inverse(Sigma.B[,])
Sigma.B[1,1] <- pow(sigma.int,2)
sigma.int ~ dunif(0, 100)          # SD of intercepts
Sigma.B[2,2] <- pow(sigma.betal,2)
sigma.betal ~ dunif(0, 100)        # SD of slopes
Sigma.B[1,2] <- rho*sigma.int*sigma.betal
Sigma.B[2,1] <- Sigma.B[1,2]
rho ~ dunif(-1,1)
covariance <- Sigma.B[1,2]

beta2 ~ dnorm(0, 0.001)
beta3 ~ dnorm(0, 0.001)
beta4 ~ dnorm(0, 0.001)
tau <- 1 / ( sigma * sigma)        # Residual
sigma ~ dunif(0, 100)              # Residual standard deviation

##priors for shade (suffix=S)
alphaS ~ dnorm (0, 0.001)          #shade intercept
beta1S ~ dnorm(0, 0.001)
beta2S ~ dnorm(0, 0.001)
beta3S ~ dnorm(0, 0.001)
idels ~ dunif(0, 100)              # Residual random variable

for (i in 1:n.sites)
{
  for (k in 1:2)
    {tauSh[i,k]<-1/Logit_shade_var[i,k] #Fixed shade variance assignment by
                                         # pre/post harvest
    }
}

delta_sim<-mean(delta_muT_sim)### Mean of the simulated temperature increase
                                   ### across sites

##### Likelihood #####
for (i in 1:n.sites)

```

```

{
  muS[i,1]<-LogitShade[i,1] #Pre-harvest value for shade uses measured
                           #values
  InvLogitMuS[i,1]<-(1/(exp(-(muS[i,1]+2.144418834))+1)) - 0.882449112
## Inverse of the logit. Both the logit of shade values and shade values
## are centered. The 2.144 value un-centers the logit of shade values.
## The 0.8824 value centers the resulting shade (between 0 and 1) value.

# Post-harvest shade estimation #
Sh_post[i,2] ~ dnorm(LogitShade[i,2],idelS) #idelS= The 'residual' random
        # variable, process variance (idelSh = i-delta-shade)
LogitShade[i,2] ~ dnorm(muS[i,2], tauSh[i,2] ) #Logit of shade values (real)
# modeled as having a mean of MuS and fixed sampling variance tauSh (real,
# see above)

muS[i,2]<-alphaS + beta1S*PctDiffBAlt100[i] + beta2S*Pre100PctHWD[i] +
beta3S*TreeHt100m[i] #MuS = equation describing mean response

InvLogitMuS[i,2]<-(1/(exp(-(muS[i,2]+1.774318915))+1)) - 0.833159815
  # needed to alter the estimated shade value, this time for post-
  # harvest. Numeric values differ because centering values differed
  # between pre-harvest and post-harvest

#change in temperature between the estimated temp and the estimated temp with
no harvest effect (NE)

delta_muT_sim[i]<- estTemp[i]- estTempNE[i] #site-level predicted increase in
                                             #temperature. See below.

#estimated temperatures from simulated harvest

  muS_calc[i]<-alphaS + beta1S*PctDiffBASim[i] + beta2S*Pre100PctHWD[i] +
beta3S*TreeHt100m[i] #estimated shade value given PctDiffBASim and other
                    # variables

  InvLogitMuS_calc[i]<-(1/(exp(-(muS_calc[i]+1.774318915))+1)) - 0.833159815

#now sticking the estimated shade into the temperature model
  estTemp[i]<-alpha[site[i]] + beta1[site[i]]*C_temp[i,1,2] +
beta2*TReachLength[i] +
  beta3*(InvLogitMuS_calc[i]) + beta4*Grad_1stQuart[i]

#estimated temperatures from no harvest

#estTempNE is for the estimated tempature, no harvest effect. The harvest
# effect is replaced by zero minus the centering amount for PctDiffBAlt100,
# 0.193039105

#Estimated NE temp shade transformation
muS_calcNE[i]<-alphaS + beta1S*(-0.193039105) + beta2S*Pre100PctHWD[i] +
beta3S*TreeHt100m[i]

```

```

InvLogitMuS_calcNE[i]<-(1/(exp(-(muS_calcNE[i]+1.774318915))+1)) -
0.833159815

#now sticking the estimated no-effects shade into the temperature model
estTempNE[i]<-alpha[site[i]] + beta1[site[i]]*C_temp[i,1,2] +
beta2*TReachLength[i] + beta3*(InvLogitMuS_calcNE[i]) +
beta4*Grad_1stQuart[i]

### Estimating temperature parameters ###
for(k in 1:2)
{
  for(j in 1:n.instances[i,k])
  {
    T_temp[i,j,k] ~ dnorm(mu[i,j,k], tau)      # The 'residual' random variable

    mu[i,j,k] <- alpha[site[i]] + beta1[site[i]]*C_temp[i,j,k] +
    beta2*TReachLength[i] + beta3*InvLogitMuS[i,k] + beta4*Grad_1stQuart[i]

  }
}

## parameter values for reporting
jags.params <- c("beta2", "beta3", "beta4", "mu.int", "sigma.int",
"mu.betal", "sigma.betal", "rho", "covariance", "sigma", "alphaS", "beta1S",
"beta2S", "beta3S", "idelS", "delta_muT_sim", "delta_sim")

## initial values
jags.inits<- function(){ list(mu.int = rnorm(1, 0, 1), sigma.int = rlnorm(1),
mu.betal = rnorm(1, 0, 1), sigma.betal = rlnorm(1), rho = runif(1, -1, 1),
sigma = rlnorm(1), beta2=rnorm(1), beta3=rnorm(1), beta4=rnorm(1), alphaS =
rnorm(1), beta1S = rnorm(1), beta2S=rnorm(1), beta3S=rnorm(1), idelS =
rlnorm(1))}

```

A2.4 As-harvested Parameter Estimates

Below are listed the parameter estimates from the As-harvested model presented in A2.3. Individual random effects estimates from the temperature sub-model are listed as $\alpha[]$ and $\beta[]$. All parameters are alphabetically listed.

	mean	sd	2.50%	25%	50%	75%	97.50%
alpha[1]	0.549	0.374	-0.197	0.291	0.562	0.779	1.326
alpha[2]	0.693	0.199	0.303	0.556	0.695	0.825	1.088
alpha[3]	1.234	0.278	0.696	1.048	1.224	1.441	1.772
alpha[4]	1.427	0.278	0.867	1.241	1.427	1.618	1.970
alpha[5]	0.419	0.314	-0.208	0.209	0.402	0.627	1.032
alpha[6]	0.101	0.256	-0.429	-0.056	0.102	0.272	0.579
alpha[7]	0.664	0.235	0.194	0.503	0.668	0.816	1.110
alpha[8]	1.084	0.214	0.675	0.939	1.085	1.235	1.496
alpha[9]	-0.341	0.377	-1.052	-0.613	-0.345	-0.069	0.342
alpha[10]	0.104	0.216	-0.308	-0.053	0.097	0.255	0.534

alpha[11]	0.825	0.257	0.342	0.649	0.817	0.994	1.314
alpha[12]	-1.880	0.199	-2.251	-2.018	-1.885	-1.743	-1.471
alpha[13]	0.267	0.197	-0.107	0.136	0.264	0.403	0.663
alpha[14]	0.460	0.312	-0.191	0.261	0.460	0.663	1.060
alpha[15]	0.816	0.199	0.448	0.674	0.810	0.949	1.200
alpha[16]	-0.141	0.346	-0.841	-0.365	-0.123	0.088	0.485
alpha[17]	0.529	0.238	0.080	0.374	0.521	0.674	1.015
alpha[18]	0.665	0.388	-0.089	0.414	0.660	0.904	1.417
alpha[19]	-0.204	0.542	-1.291	-0.561	-0.210	0.138	0.916
alpha[20]	0.016	0.207	-0.417	-0.119	0.024	0.152	0.402
alpha[21]	0.667	0.214	0.244	0.524	0.673	0.820	1.085
alpha[22]	-0.216	0.487	-1.173	-0.535	-0.202	0.103	0.747
alpha[23]	0.354	0.323	-0.264	0.138	0.353	0.567	1.013
alpha[24]	-0.002	0.328	-0.637	-0.231	0.004	0.217	0.642
alpha[25]	0.233	0.413	-0.559	-0.022	0.223	0.482	1.062
alpha[26]	1.162	0.262	0.658	0.986	1.155	1.335	1.679
alpha[27]	0.871	0.353	0.172	0.634	0.860	1.104	1.559
alpha[28]	1.125	0.222	0.697	0.978	1.125	1.271	1.568
alpha[29]	-0.304	0.383	-1.091	-0.565	-0.289	-0.037	0.425
alpha[30]	-0.016	0.279	-0.567	-0.196	-0.014	0.167	0.544
alpha[31]	0.861	0.264	0.340	0.683	0.855	1.049	1.377
alpha[32]	0.677	0.264	0.130	0.506	0.682	0.843	1.178
alpha[33]	0.385	0.267	-0.192	0.216	0.395	0.566	0.894
alphaS	-0.279	0.066	-0.407	-0.321	-0.279	-0.237	-0.148
beta1[1]	-0.380	1.839	-3.681	-1.605	-0.486	0.769	3.655
beta1[2]	-0.750	1.742	-4.098	-1.928	-0.803	0.403	2.756
beta1[3]	-2.908	0.579	-4.039	-3.300	-2.912	-2.528	-1.805
beta1[4]	0.044	0.719	-1.428	-0.426	0.065	0.540	1.342
beta1[5]	-0.015	1.403	-2.595	-0.974	-0.100	0.906	3.010
beta1[6]	-0.651	1.726	-3.916	-1.852	-0.677	0.468	2.779
beta1[7]	-0.499	0.439	-1.327	-0.794	-0.512	-0.190	0.389
beta1[8]	-1.268	1.799	-4.727	-2.428	-1.217	-0.089	2.430
beta1[9]	-2.129	0.599	-3.310	-2.532	-2.116	-1.716	-0.985
beta1[10]	-0.656	1.933	-4.170	-1.962	-0.779	0.553	3.490
beta1[11]	0.714	1.840	-2.583	-0.508	0.611	1.899	4.707
beta1[12]	-2.806	1.026	-4.938	-3.451	-2.828	-2.156	-0.688
beta1[13]	-3.116	0.475	-4.063	-3.435	-3.134	-2.799	-2.141
beta1[14]	-3.602	0.907	-5.355	-4.185	-3.580	-2.961	-1.907
beta1[15]	-0.933	0.905	-2.617	-1.573	-0.938	-0.315	0.902
beta1[16]	-0.993	1.722	-4.316	-2.104	-1.075	0.094	2.378
beta1[17]	-1.088	1.660	-4.385	-2.166	-1.088	0.109	2.090
beta1[18]	-1.532	0.714	-3.041	-1.985	-1.509	-1.094	-0.191
beta1[19]	-0.206	1.219	-2.724	-0.985	-0.227	0.535	2.282
beta1[20]	-0.110	1.703	-3.161	-1.233	-0.249	0.908	3.578
beta1[21]	-0.029	0.935	-1.912	-0.649	-0.035	0.610	1.786

beta1[22]	-2.450	0.845	-4.140	-3.010	-2.412	-1.883	-0.826
beta1[23]	-2.396	1.366	-5.153	-3.249	-2.330	-1.521	0.335
beta1[24]	-1.230	0.843	-2.818	-1.813	-1.243	-0.664	0.467
beta1[25]	-1.332	1.336	-4.008	-2.161	-1.365	-0.479	1.374
beta1[26]	0.453	1.135	-1.876	-0.309	0.502	1.231	2.612
beta1[27]	-1.030	0.665	-2.344	-1.462	-1.012	-0.589	0.234
beta1[28]	-3.332	0.609	-4.525	-3.717	-3.355	-2.937	-2.132
beta1[29]	-1.962	1.850	-5.665	-3.159	-1.949	-0.794	1.630
beta1[30]	-0.462	1.396	-3.087	-1.409	-0.501	0.447	2.317
beta1[31]	0.109	1.778	-3.150	-1.153	0.042	1.203	3.867
beta1[32]	-1.043	1.737	-4.561	-2.194	-1.022	0.025	2.431
beta1[33]	2.312	1.250	-0.064	1.493	2.286	3.162	4.711
beta1S	-2.776	0.305	-3.428	-2.973	-2.770	-2.558	-2.223
beta2	0.871	0.336	0.206	0.663	0.878	1.077	1.511
beta2S	-0.585	0.249	-1.092	-0.754	-0.583	-0.414	-0.100
beta3	-5.606	0.844	-7.341	-6.153	-5.590	-5.030	-4.046
beta3S	-0.065	0.017	-0.100	-0.076	-0.066	-0.054	-0.031
beta4	-0.077	0.049	-0.179	-0.109	-0.076	-0.044	0.014
covariance	0.218	0.351	-0.435	-0.011	0.211	0.429	0.954
delta_muT_sim[5]	0.300	0.045	0.219	0.267	0.300	0.330	0.392
delta_muT_sim[8]	1.069	0.116	0.855	0.984	1.068	1.148	1.291
delta_muT_sim[9]	1.309	0.148	1.033	1.200	1.308	1.414	1.594
delta_muT_sim[10]	2.109	0.209	1.715	1.967	2.108	2.255	2.499
delta_muT_sim[11]	1.152	0.113	0.923	1.074	1.153	1.227	1.370
delta_muT_sim[12]	0.448	0.051	0.354	0.411	0.448	0.484	0.544
delta_muT_sim[15]	-0.066	0.007	-0.081	-0.071	-0.066	-0.061	-0.052
delta_muT_sim[18]	0.303	0.055	0.207	0.264	0.300	0.339	0.421
delta_muT_sim[19]	0.306	0.044	0.222	0.276	0.306	0.335	0.394
delta_muT_sim[20]	1.134	0.136	0.856	1.044	1.133	1.227	1.394
delta_muT_sim[21]	0.974	0.095	0.785	0.909	0.975	1.037	1.151
delta_muT_sim[22]	1.331	0.163	1.006	1.218	1.332	1.438	1.652
delta_muT_sim[23]	0.358	0.041	0.278	0.330	0.358	0.385	0.439
delta_muT_sim[24]	2.196	0.219	1.769	2.048	2.195	2.346	2.606
delta_muT_sim[25]	2.173	0.213	1.763	2.026	2.172	2.312	2.569
delta_muT_sim[27]	0.212	0.023	0.169	0.196	0.211	0.227	0.257
delta_muT_sim[30]	0.591	0.072	0.457	0.539	0.590	0.639	0.731
delta_muT_sim[33]	0.913	0.102	0.717	0.843	0.910	0.981	1.114
delta_muT_sim[1]	-0.050	0.007	-0.065	-0.055	-0.050	-0.045	-0.036
delta_muT_sim[2]	-0.169	0.021	-0.209	-0.183	-0.168	-0.155	-0.130
delta_muT_sim[3]	-0.020	0.003	-0.026	-0.022	-0.020	-0.018	-0.015
delta_muT_sim[4]	0.036	0.004	0.028	0.033	0.035	0.038	0.044
delta_muT_sim[6]	-0.027	0.003	-0.034	-0.030	-0.027	-0.025	-0.021
delta_muT_sim[7]	-0.003	0.000	-0.004	-0.003	-0.003	-0.003	-0.002
delta_muT_sim[13]	-0.006	0.001	-0.008	-0.007	-0.006	-0.006	-0.005
delta_muT_sim[14]	0.116	0.012	0.093	0.107	0.115	0.125	0.139

delta_muT_sim[16]	0.089	0.011	0.067	0.081	0.090	0.097	0.112
delta_muT_sim[17]	-0.098	0.014	-0.124	-0.107	-0.097	-0.089	-0.072
delta_muT_sim[26]	0.054	0.006	0.043	0.050	0.054	0.058	0.065
delta_muT_sim[28]	-0.154	0.034	-0.226	-0.175	-0.151	-0.128	-0.097
delta_muT_sim[29]	-0.072	0.011	-0.094	-0.080	-0.072	-0.065	-0.053
delta_muT_sim[31]	-0.465	0.069	-0.602	-0.513	-0.464	-0.419	-0.327
delta_muT_sim[32]	-0.017	0.002	-0.021	-0.019	-0.017	-0.016	-0.013
delta_sim	0.486	0.049	0.395	0.451	0.485	0.519	0.576
deviance	85.727	17.489	54.325	73.848	84.686	96.386	122.888
idels	49.967	28.579	2.444	25.738	48.646	75.080	96.816
mu.beta1	-1.092	0.460	-1.951	-1.399	-1.105	-0.807	-0.152
mu.int	0.396	0.133	0.104	0.310	0.400	0.484	0.645
rho	0.153	0.228	-0.297	-0.007	0.165	0.319	0.555
sigma	0.280	0.029	0.233	0.260	0.277	0.299	0.345
sigma.beta1	1.921	0.440	1.186	1.599	1.887	2.199	2.861
sigma.int	0.730	0.111	0.538	0.654	0.720	0.794	0.981

Appendix 3 - State forest harvest, as programmed

The state forest harvest treats stand data similarly to the FPA harvest (Appendix 4) except that tree distances are measured as horizontal distance from the stream. Within this program, trees are preferentially retained closest to the streams within RMA zones. The structure of the harvest simulation differs from the FPA harvest in that the inner zone is first assessed to determine if the stand meets mature forest condition (MFC). Then the stand data are passed on to functions to simulate harvest in the inner zone and the outer zone (did not meet MFC) or just the outer zone (MFC achieved).

The approach for this simulated harvest was gleaned from Appendix J and Appendix C of the Northwest Oregon Forest Management Plan as well as from communications with State Forest employees. RipStream vegetation data are well-suited for applying Appendix J to, although there are some data limitations. Limitations, assumptions, and decisions include:

- Only have data for trees >6" DBH
- Only have crown class/height information for roughly 1/5 of trees
- Will determine SDI only for conifer species >8" DBH, see table J-1 in NWFMP. This diameter best matches my understanding of how field offices calculate SDI for their stands.
- All hardwoods within 100' of streams in all stands will be retained.
- SDI is a mixed-species approach, and I sum % SDI by conifer species
- There are two conifer MFC criteria, it appears. The first is that there needs to be 220 square feet of 11" conifer basal area/acre in the inner zone. Later on, when figuring out how many trees to retain in the outer RMA zone, it defines a MFC of having >45 11" conifer TPA in the inner zone (first 100'). 45 11" trees = 119 square feet. Therefore it does not appear that the two measures are equivalent. The first value is calculated to determine if the inner zone receives harvest. Next, regardless of selected harvest, the second MFC (>45 11" conifer/ac) is applied to determine outer zone retention.
- In outer zone, we select conifers for retention based on size, in order from largest to smallest. Regardless of distance from the stream, the largest are retained (between 8 & 35 per 500')
- If MFC is attained because the stand is hardwood dominated, then the outer zone has 8 conifers retained + the deficit number. If MFC is achieved with large conifers, then the outer zone MFC is set to 8 (it would be 15 if 1000' of stream).
- When calculating SDI, we count unknown conifer species as "DF."

Table J-1. Management Standards for Type F Stream RMAs

All Stream Sizes: Large, Medium, and Small	
Stream bank zone 0-25 ft.	<ul style="list-style-type: none"> • No harvest. • Less than 10% vegetative disturbance. • Full suspension required during cable yarding. • No ground-based equipment operation. • Leave any trees damaged or felled from yarding activities.
Inner RMA zone 25 to 100 ft.	<ul style="list-style-type: none"> • Manage for mature forest condition.¹ • No management activity where mature forest condition (MFC) exists, or where conditions are suitable for development of MFC in a reasonable time frame without further treatment. • Actively manage where necessary to achieve the desired future condition in a timely manner. • Minimum 15-year interval between harvest entries, and minimum number of entries necessary to achieve the desired future condition. • Partial cutting will maintain a conifer density of at least SDI 25%, and will retain at least 50 TPA. • No more than 10% vegetative disturbance allowed from cable yarding. • Full suspension wherever possible, or one-end suspension on all cable-yarded material. • Ground-based equipment operation limited to area more than 50 ft. from aquatic zone and slopes less than 35%, and allowed on no more than 10% of area. • Leave any trees damaged or felled from yarding activities and additional felled, girdled or topped trees to contribute toward down wood targets.² • Retain all dead and down material that was present prior to the operation.
Outer RMA zone 100 to 170 ft.	<ul style="list-style-type: none"> • Retain at least 10 to 45³ conifer trees and snags per acre (15 to 70 trees per 1,000 ft. of RMA).⁴ • Retain all snags as safety permits. • Less than 10% ground disturbance from yarding activities. • Retain all dead and down material that was present prior to the operation.

1. Desired mature forest condition consists of a stand dominated by large conifer trees, or where hardwood-dominated conditions are expected to be the natural plant community, a mature hardwood/shrub community. For conifer stands, this equates to a basal area of 220 square feet or more per acre, inclusive of all conifers over 11 inches DBH. At a mature age (80-100 years or greater), this equals 40-45 conifer trees 32 inches in DBH per acre.
2. Up to 10 trees per acre will be retained as felled, girdled, or topped trees during partial cutting, to reach a target of 600-900 cubic feet per acre of hard down wood.
3. Outer zone tree retention target will be increased when less than the target number of conifers is present in the inner zone. The process for calculating the outer zone retention target is described in the section following the RMA prescription tables.
4. All trees retained will be dominant or co-dominant conifer trees (if available). In order to balance the need for short-term and long-term recruitment of large wood to the aquatic zone, preference will be given to retaining trees on adjacent slopes, trees leaning toward the aquatic zone, and trees closest to the channel.

Table J1 from the NWFMP

From Appendix J, page 11, NWFMP:

Increasing Outer Zone Conifer Retention on Type F Streams

On Type F streams, in situations where the number of conifers available for retention within the inner zone is not adequate to achieve the large wood delivery potential of a mature forest condition, additional conifers will be retained in the outer zone to provide additional large wood recruitment potential.

This additional outer zone target will apply when the number of conifers of suitable size (11 inches or greater DBH) in the inner zone is less than the mature forest condition target of 45 TPA (100 trees per 1,000 lineal feet of stream for a 100-foot inner zone).

The number of additional conifers to be retained in the outer zone will be equal to the deficit from the inner zone target, adjusted to account for the different widths of the zones. For example, if the inner zone has an average of 70 suitable conifers per 1,000 feet of stream, then the additional retention level for the outer zone would equal 30 times 0.7, or an additional 21 conifers per 1,000 feet of outer zone.

In no case shall the number of conifers required to be retained in the outer zone exceed the inner zone target for mature forest condition. This means no more than 70 conifers per 1,000 feet of outer zone or 45 TPA are required. In addition, no trees shall be required to be retained in the outer zone in locations where, due to topography, they would have no opportunity to reach the area within the channel migration zone and thus potentially function as large wood in the stream channel. All conifers retained under this strategy shall meet the conifer retention criteria as described in footnotes to Tables J-1 and J-2: dominant or co-dominant trees, with preference given to retaining trees on adjacent slopes, trees leaning toward the aquatic zone, and trees closest to the channel.

Following table J1 and instructions on increasing outer-zone conifer retention on F streams, we altered the pre-harvest stand data according to the following:

Data Preparation

Pre-harvest vegetation plot data are first cleaned to make sure only trees with DBH ≥ 6 " and with horizontal distances ≤ 170 are included. Then the data are reduced to a single plot. All trees are given a rank by distance. Distances are minutely and randomly adjusted to break ties.

The next step is to determine if stands are at MFC (either hardwood or conifer) or not. Conifer MFC is relatively easy to determine. It is the sum of all living conifer basal area between 25 & 100' from the stream (inner zone) with a DBH ≥ 11 ". This amount is converted to an acreage value and compared against the target (220 ft²/ac).

The data are also passed to the function HW_DOMINATION to determine if hardwood dominated (described below). If the plot is either conifer or hardwood dominated, the tree data are passed to the function MFC_PATH in which no harvest happens in the inner zone. If not hardwood or conifer dominated, then data are passed to the function SDI_PATH, which determines harvest levels in the inner and outer zone.

HW_DOMINATION

The following description of hardwood stands is found in Appendix C, page 15, of the NWFMP.

Hardwoods

Hardwood stands are classified along with conifer stands in one of the five stand structure types. However, for the purpose of facilitating discussion, hardwood stands are defined as those stands where hardwood tree species comprise more than 70 percent of the tree canopy. Seventy percent is a subjectively set measure that identifies when the hardwood canopy is the dominant vegetative feature that characterizes the stand tree canopy and thus will likely control the focus of stand management practices. Seventy percent is also being used to identify hardwood stands by current research such as the “Coastal Landscape Analysis and Modeling Study” (CLAMS) (Tom Spies 1996). Common hardwood tree species include red alder, bigleaf maple, and Oregon white oak.

Field managers may choose to manage hardwood stands on the landscapes for a variety of reasons, such as to obtain economic benefits from hardwood products, to manage tree diseases in the stand, or to introduce or maintain additional vegetative diversity within conifer-dominated landscapes.

At this time it is assumed that a small percentage (probably 10 percent or less) of the landscape will be managed as hardwood stands. Maintaining a component of hardwoods within conifer stands is encouraged and it is anticipated that most stands will have some hardwoods. Implementation plans will better estimate how much of the landscape currently consists of hardwood stands and what portions of the landscape may be managed as hardwood stands in the future. If managers determine it is desirable to manage greater portions of the landscape in hardwoods, the forest management plan may have to be adjusted.

To determine if a plot is hardwood dominated, only the live trees in the inner zone (25-100') are considered. Hardwood dominated was interpreted as > 70% of the canopy. However, we do not have full canopy class (CC) information for our stands. Below is how we interpreted our stand data to determine hardwood dominance.

For inner zone, remove all snags, and create list of all trees with crown class of 4+ (intermediate, overtopped) and another list of trees with CC= 2, 3 (dominant, codominant).

For just the tree species that had dominant, codominant stems, stems were treated as canopy stems if DBH was greater than the maximum stem DBH for intermediate/overtopped trees of that species. Or, if all measured stems were in the canopy, we included all stems that were greater than the minimum-sized CC 2,3 stem for that species.

For all identified “canopy” stems we determined the fraction of those stems that were hardwood. If that fraction was ≥ 0.7 the stand was coded as hardwood-dominant (1) and therefore MFC was reached. If not, it was coded a zero (0). This value was passed back to the Data Preparation step to contribute towards determining the destination of the data (functions MFC_PATH or SDI_PATH).

MFC_PATH

If the site’s inner RMA zone is at MFC

Keep all trees $\leq 100'$ from stream

Determine if secondary conifer MFC condition met (100 trees/1000' DBH $\geq 11''$)

- a. If condition met, the 8 largest (DBH) conifer trees in outer zone are retained, all else harvested
- b. If condition not met, determine the number of trees short of the secondary MFC. The number must be ≤ 35 trees/500' and 8 or greater.

Send plot information to the **Plot Summarization** function

SDI_PATH

If RMA zone not at MFC, then determine amount of conifer to be harvested according to SDI.

Compile a stand list of conifers (one plot). The list is sorted by distance from stream.

The inner zone is defined as $>25, <101$ feet

For each tree over 11" DBH, the SDI is calculated

- a. A blank table is created with columns = conifer tree species
- b. Cycling through each tree in the inner zone, the SDI table is updated. The number of trees/tpa for that species is increased, as is the cumulative basal area for that species. These numbers are used to calculate the QMD and SDI for that species. Then, the SDI is compared against the maximum SDI for that species to produce the percentage SDI achieved by that species. The percentages for all tree species are totaled and recorded as a cumulative value for the target tree.
- c. $QMD = \sqrt{BA/(k*n)}$, where BA = stand basal area, k in square feet = 0.005454, and n = number of trees
- d. $SDI = tpa*(QMD/10)^{1.605}$

Trees are marked in the inner zone as needing to be kept (in order of distance from the stream) if the $SDI < 0.25$ and there are less than 50 tpa (59 trees for a 500'*75' plot). The program is set to keep all 8+'' trees that contribute to SDI, selected in order by distance from the stream. Trees contributing to the SDI are spared, as are additional trees needed to reach the 50 tpa target. If the SDI is met but fewer than 50 tpa have been retained, 8+'' trees are retained preferentially over smaller trees. If smaller trees are needed to meet the 50 tpa target, they will be retained starting from the 25' distance. All other conifers will be removed.

- a) The tpa measure can be met with $<8''$ trees, although this is anticipated to be rare since it would take a number of large conifers to have the SDI reach 25% but not 50 TPA
- b) Once we reach SDI, a list of trees to keep is determined
 - a. Keep all trees within 25' from stream
 - b. Keep all hardwood within 100'
 - c. Keep all SDI/tpa selected conifers from inner zone (harvest the rest)
 - d. Keep selected conifers from outer zone as per step b in the function MFC_Path.

Plot data are then sent to the **Plot Summarization** function; all trees (kept or otherwise) were passed on to the function **Tree Fate**.

Plot Summarization

As in the FPA harvest simulation, Plot Summarization takes the appropriately-annotated vegetation plot data that passed through the MFC_PATH or SDI_PATH functions and summarizes the remaining basal area. Incoming data have been stripped of snags, as snags are not expected to contribute much to shade. The stand data trees are marked as “kept” = 1 or 0, to indicate that the tree is retained or harvested. Summarized data include (within 100’ horizontal distance of the stream, for all “kept” = 1 trees) the basal area of all trees, conifer trees, hardwood trees, all trees expressed in m², the number of all trees, conifers, and hardwoods, and the distance to the retained trees farthest from the stream or (of the five vegetation plot lines) the mean and minimum of the maximum tree distance along each line.

The output file for Plot Summarization provides the information for each plot on a single line.

Tree Fate

The Tree Fate function is a means for storing all of the tree data along with a column for “kept” trees. The idea is to be able to plot the retention patterns for each plot as needed.

Appendix 4 - FPA harvest, as programmed

The FPA analysis differs from the State analysis in that it examines the numbers of trees within a specified slope distance, not horizontal distance, from the stream. It is similar in that it preferentially retains trees closest to the stream.

The annotated FPA rules for Type F streams (629-640-0100)

629-640-0100

General Vegetation Retention Prescription for Type F Streams

(1) (a) *Operators shall apply the vegetation retention requirements described in this rule to the riparian management areas of Type F streams.*

(b) *Segments of Type F streams that are different sizes within an operation shall not be combined or averaged together when applying the vegetation retention requirements.*

(c) *Trees left to meet the vegetation retention requirements for one stream type shall not count towards the requirements of another stream type.*

We applied (1) (a) to the vegetation data as a whole as closely as possible. (b) and (c) did not affect the analysis.

(2) *Operators shall retain:*

(a) *All understory vegetation within 10 feet of the high water level;*

(b) *All trees within 20 feet of the high water level; and*

(c) *All trees leaning over the channel.*

We apply (2) (b) to the analysis but not (2) (a) or (2) (c). We did have tree lean for some trees (towards or away from stream) but no information on whether the lean was over the channel.

(3) *Operators shall retain within riparian management areas and streams all downed wood and snags that are not safety or fire hazards. Snags felled for safety or fire hazard reasons shall be retained where they are felled unless used for stream improvement projects.*

Not applicable to the simulated harvest.

(4) *Notwithstanding the requirements of section (2) of this rule, vegetation, snags and trees within 20 feet of the high water level of the stream may be felled, moved or harvested as allowed in other rules for road construction, yarding corridors, temporary stream crossings, or for stream improvement.*

Not applicable to the simulated harvest.

(5) *Operators shall retain at least 40 live conifer trees per 1000 feet along large streams and 30 live conifer trees per 1000 feet along medium streams. This includes trees left to meet the requirements described in section (2) of this rule. Conifers must be at least 11 inches DBH for large streams and 8 inches DBH for medium streams to count toward these requirements.*

Applied; see description below.

629-640-0100(6) is elaborated on below. Section 629-640-0100(7) is incorporated in the analysis of 629-640-0100(6)(a). We did not apply 629-640-0100(8) through (13).

Data Preparation

The FPA harvest (and all of the harvest files) examine the pre-harvest RipStream vegetation plot cruise data.

The program at first assesses the type of riparian retention to take place at the unit. That is, it needs to determine whether the Type 3 harvest (clearcut) will be 629-640-0100(6)(a, b, or c).

I refer to these three Alternatives as Alt A, B, and C.

It sets a no-cut distance of 20' for all sites, except for medium Alt C streams.

Stream sizes are obtained from a file.

We considered only the treatment reach plots (plots 1 & 2).

All trees are removed from consideration that have $DBH < 6''$.

Finally, I reduced the data to trees from a single vegetation plot.

Determining harvest fates Alts 6a, 6b, 6c, or No Conifer Harvest

This section of the program determines a "harvest function" to pass the plot data to. The data are not reduced here, just sent on their way to their appropriate destination.

Given the stream size (medium or small) the overall file was reduced to include all trees within 70' (Medium) or 50' (Small).

Standard targets. These were set to $BA = 60'$ or $20'$ for medium or small streams in the coast range. Site 7353 was in the interior and a medium stream, so its' standard target was set to 70. Notice these targets are 50% of what they should be. This is because our plots were 500' long, not 1000'.

I then ranked all trees by distance, with 1 being closest to the stream. A random number between 0 and 1 was associated with these distances to break ties.

Next, to figure out how much snag and hardwood basal area is available & usable to substitute for conifer basal area, I determine the basal area of all conifers within the RMA, the basal area of conifer snags with heights $> 30'$, and the basal area of non-(red) alder, non-conifer, non-snag hardwoods $> 20'$ from the stream and 24" DBH or greater. I determined whether the conifer basal area within the RMA was $> \text{or} = 90\%$ of the standard target, if the available snag & hardwood basal area was $> 10\%$ of the standard target, and finally the percent that snags and hardwood contribute can contribute to the standard target (max of 10%).

The number of non-snag conifers with $DBH > 8''$ are tallied as well.

Harvest Alt 6a:

If the conifer standard target can be met or exceeded then 6a is selected.

Harvest Alt 6b:

If conifer basal area is $> 50\%$ but $< 100\%$ of the standard target then 6b is selected.

Harvest Alt 6c:

If $< 50\%$ of the standard target may be met with conifer basal area then 6c is selected.

No Conifer Harvest:

If the stream size is "medium" and there are fewer than 15 8" DBH conifers within the 500' section of RMA (30 per 1000') then no conifer harvest of the RMA is permitted.

Once a harvest alternative is selected, the stand data are passed on to the appropriate function.

The harvest functions determine which trees stay and which trees are harvested.

Harvest Alt 6a

- (6) Operators shall retain trees or snags six inches or greater DBH to meet the following requirements (this includes trees left to meet the requirements of sections (2) and (5) of this rule):
- (a) If the live conifer tree basal area in the riparian management area is greater than the standard target shown in Table 2 [on page 49] where the harvest unit will be a harvest type 2 or type 3 unit (as defined by ORS 527.620), or Table 3 [on page 50] where the harvest unit will be a harvest type 1, partial harvest, or thinning, operators shall retain live conifer trees of sufficient basal area to meet the standard target.
 - (7) In the Coast Range, South Coast, Interior, Western Cascade, and Siskiyou geographic regions, hardwood trees and snags six inches or greater DBH may count toward the basal area requirements in subsection (6)(a) of this rule as follows:
 - (a) All cottonwood and Oregon ash trees within riparian management areas that are beyond 20 feet of the high water level of large Type F streams, may count toward the basal area requirements.
 - (b) Up to 10 percent of the basal area requirement may be comprised of sound conifer snags at least 30 feet tall and other large live hardwood trees, except red alder, growing in the riparian management area more than 20 feet from the high water level and at least 24 inches DBH.

We attempted to program the harvest procedure to be as true to (6) (a) as possible and to incorporate available basal area according to (7). The code involved largely attempted to determine which trees were necessary to meet the 30 per 1000' of stream (for Medium streams) and if the standard target was met, using as much basal area credit as was possible from conifer snags and large non-alder hardwoods >20' from the stream's edge and 24" DBH or greater . We assumed a Type 3 harvest (clearcut) would take place.

The first task is to determine 30 live conifers per 1000' of medium stream. The distance ranks from above are used to rank all 8" or greater conifers. Also, the number of these types of conifers is recorded.

Next, the code selects which conifer snags or large hardwood trees to retain and count towards the standard target. All possible trees are identified, and then they are ranked by distance from stream and their total basal area contribution summed. Then the basal area of these hardwoods and snags is summed for each of these trees and all qualifying snags/hardwoods closer to the stream. The program finds out the minimum basal area sum that is sufficient to meet or exceed 10% of the standard target (in this case, 2 or 6 square feet of BA for small and medium streams, respectively. $2 \times 2 = 4 = 10\%$ of 40' for small streams; $6 \times 2 = 12$ or 10% of 120'). If there is < 10% of the standard target that can be met by conifer snags and hardwoods, then the basal area of all contributing hardwoods and snags is recorded. The code also records the amount by which the standard target may be reduced for conifer.

Similar to the conifer snags and hardwood, conifers are ranked by distance from the stream and then their cumulative basal area summed, by tree. The tree with the cumulative basal area value (to and including that tree) greater than the standard target minus the snag/hardwood reduction amount is identified.

Next we go to the tree selection process.

Conifer trees are not harvested if:

- It is one of the 15 8" conifers necessary to meet stem number requirements along "medium" streams (30/1000)
- It is one of the conifer snags or large hardwoods identified to meet the standard target amount
- It is one of the identified conifers necessary to meet the standard target

- The tree is \leq the no-cut width.

All hardwoods to 20' of the channel are removed.

Of note, aside from particularly large non-alder hardwoods $> 20'$ from the channel and all hardwoods within 20' of the channel, ALL hardwoods are harvested from the RMA. This is not true if there are fewer than 15 conifers $> 8''$ DBH (which would have triggered the No Conifer Harvest prescription).

The plot data at this point has some extra columns, that identify which trees are to be kept and which are not. This information is sent on to the **plot summarization** function and to the **tree fate** function. We re-insert all snags into the tree fate data set, but keep the snags out of the plot summarization data set.

Harvest Alt 6b

- (6) Operators shall retain trees or snags six inches or greater DBH to meet the following requirements (this includes trees left to meet the requirements of sections (2) and (5) of this rule):
- (b) If the live conifer tree basal area in the riparian management area is less than the standard target (as shown in Table 2 where the harvest unit will be a harvest type 2 or type 3 unit, or Table 3 where the harvest unit will be a harvest type 1, partial harvest, or thinning) but greater than one-half the standard target shown in Table 2, operators shall retain all live conifer trees six inches DBH or larger in the riparian management area (up to a maximum of 150 conifers per 1000 feet along large streams, 100 conifers per 1000 feet along medium streams, and 70 conifers per 1000 feet along small streams).

Harvest 6b differs from 6a in that snags and hardwoods may not count towards the standard target, as the standard target will by definition not be met. Similar to 6a, 30 conifers $> 8''$ DBH are needed per 1000' of stream channel, so the program ranks 8'' conifers by their distance to the stream and determines how many there are. It ranks 6'' and greater conifers by their distance to the stream, and determines their number. Conifers that may be harvested in excess of 150/1000' for Medium streams and 70/1000' for Small stream. All hardwoods to 20' of the channel are removed.

The program:

First, all non-snag conifers $\geq 8''$ DBH are counted and ranked, as they were in 6a. All conifers 6'' DBH and greater are counted and ranked as well (a different ranking system than the 8'' + trees).

Tree selection process:

- All trees within the no-cut distance are retained
- If a stream is a Medium, the 15 closest conifers with a DBH $\geq 8''$ are kept
- If a stream is a Small, the 35 closest 6+'' DBH trees are kept; if the stream is Medium the tree number is increased to 50

Similar to 6a, the altered plot data are sent on to the **plot summarization** function and to the **tree fate** function. We re-insert all snags into the tree fate data set, but keep the snags out of the plot summarization data set.

Harvest Alt 6c

(6) Operators shall retain trees or snags six inches or greater DBH to meet the following requirements (this includes trees left to meet the requirements of sections (2) and (5) of this rule):

(c) If live conifer tree basal area in the riparian management area is less than one-half the standard target shown in Table 2:

- (A) Operators may apply an alternative vegetation retention prescription as described in OAR 629-640-0300 where applicable, or develop a site specific vegetation retention prescription as described in OAR 629-640-0400; or
- (B) Operators shall retain all conifers in the riparian management area and all hardwoods within 50 feet of the high water level for large streams, within 30 feet of the high water level for medium streams, and within 20 feet of the high water level for small streams.

(6)(c)(A) which provides prescriptions for catastrophic stand damage events and conversion of hardwood dominated RMA is not applicable to the simulated harvest.

For plot data passed on to the function for 6c(B), all conifers within the RMA are retained (we started with trees of 6" DBH or greater, so our data and the criteria above coincide). We use two no-cut distances in this function, set to 20' for Small streams or 30' for Medium streams. All hardwoods are harvested beyond these points by the program.

Tree selection process:

- All trees within the no-cut distances (20' or 30') are retained
- All conifers 6+" DBH are retained within the RMA

Similar to 6a, the altered plot data are sent on to the **plot summarization** function and to the **tree fate** function. We re-insert all snags into the tree fate data set, but keep the snags out of the plot summarization data set.

No Conifer Harvest

If there are fewer than 15 8+" conifers in the RMA then conifers in the entire RMA remain unharvested. All trees and snags of all species are kept and the plot data are passed on to plot summarization and tree fate. Snags were removed from the basal area of stand data passed to **plot summarization**; all trees were passed on to the function **tree fate**. Note that harvest did occur beyond 70' for medium streams and 50' for small streams.

Plot Summarization

The Plot Summarization takes the appropriately-annotated vegetation plot data (that has now passed through Harvest Type 6a, 6b, 6c, or No Conifer Harvest) and summarizes the remaining basal area. As noted above, the incoming data have been stripped of snags, as snags are not expected to contribute much to shade. The stand data trees are marked as "kept" = 1 or 0, to indicate that the tree is retained or harvested. Summarized data include (within 100' slope distance of the stream, for all "kept" = 1 trees) the basal area of all trees, conifer trees, hardwood trees, all trees expressed in m², the number of all trees, conifers, and hardwoods, and the distance to the retained trees farthest from the stream or (of the five vegetation plot lines) the mean and minimum of the maximum tree distance along each line.

The output file for Plot Summarization provides the information for each plot on a single line.

Tree Fate

The Tree Fate function is a means for storing all of the tree data along with a column for “kept” trees. The idea is to be able to plot the retention patterns for each plot as needed.

RipStream Predictive Harvest Effects Analysis

Overview & ODF analysis approach

Part I: Predictive Shade & Temperature Model

1.1 Introduction the Predictive Shade & Temperature Model

1.2 RipStream study design background

1.3 PCW analysis

1.4 Effects analysis

1.5 Need for a more flexible approach – Bayesian modeling

1.6 Developing sub-models

1.6.1 Temperature

1.6.2 Shade

1.7 Combining the sub-models

1.8 Model evaluation

1.9 Model – specific assumptions

1.10 Model predictions

Part II: Harvest Simulation Approach

2.1 Vegetation data Use

2.2 Predictions

2.2.1 As-harvested

2.2.2 State Forests

2.2.3 Private Forests

2.2.4 Percent Basal Area

2.2.5 Harvest by distance from Stream

Appendix 1: Comparing 7DMM and 40-day max values

Appendix 2: Bayesian model parameterization

Appendix 3: State forest harvest simulation

Appendix 4: Private forest harvest simulation

DRAFT

Overview

In 1994 the Oregon Department of Forestry altered its Water Classification & Protection Rules (ODF 1994) to create the current suite of Oregon Forest Practices Act (FPA) riparian rules. In 1998 a MOU was signed between the Oregon Department of Environmental Quality (DEQ) and the Oregon Department of Forestry (ODF) that agreed to provide regulatory certainty for forest practices on non-federal land under the Clean Water Act so long as ODF monitoring verified that forest practices effectively protected water quality. ODF followed through on the agreement by enacting, in 2002, the Riparian Function and Stream Temperature project (RipStream). RipStream represented a highly controlled study of 33 sites in the Oregon Coast Range on state and privately-owned land. Data collection occurred generally for two years pre-harvest and five years post-harvest at all sites. Every site had a control reach upstream of the treatment reach. Data were collected repeatedly on stream temperatures, site vegetation, channel characteristics, and shade. Three peer-reviewed publications have come out of the study. Dent et al. (2008) provided an examination of site characteristics pre-harvest. The first post-harvest analysis, Groom et al. (2011a), asked whether the DEQ Protecting Cold Water criterion was met. It appeared that that PCW was not met on privately-owned forestland. The result warranted a closer examination of the data, resulting in a second post-harvest analysis (Groom et al. 2011b). This second post-harvest analysis had three main findings: 1) Streams on private lands appeared to be warmed by 0.7 °C while State forest streams did not change in temperature on average (0.0 °C); 2) Stream temperature changes were driven by changes in shade; 3) Changes in shade were largely related to changes in basal area. These results were presented to the Oregon Board of Forestry, the political body that oversees ODF and the administration of the FPA. The Board examined the findings and concluded that degradation of cold water had occurred on private lands. This finding triggered Section 527.714 of the Forest Practices Act, which initiates an examination of rule sufficiency and consideration of rule alteration. This paper presents the methods developed to produce recommendations to the Board of Forestry on alternative rule changes necessary to protect Cold Water.

The prediction method described here is intended for developing harvest prescriptions to the Board of Forestry. Although the sites were not randomly selected (see Groom et al. 2011b), virtually all available and suitable sites were used. The purpose of this effort is to model the temperature and shade responses to harvest at all 33 sites, and use the relationships found to predict temperature changes at those 33 sites given different harvest prescriptions. The overall temperature signal due to harvests across these sites will be interpreted as representative of harvests conducted across the geographical area of interest. The model is not intended for use beyond this effort; obtaining the necessary information, such as pre-harvest and upstream control temperature behavior, will generally not be possible in other settings.

Part I: Predictive Shade & Temperature Model

1.1 Introduction to the Predictive Shade & Temperature Model

The Predictive Shade & Temperature Model (PSTM) development depended on the earlier modeling efforts, yet it differs from them in some critical ways. The Protecting Cold Water analysis (PCW analysis; Groom et al. 2011a) produced findings that spurred subsequent analyses, yet it used temperature metrics that the others did not. The effects and magnitude analysis (effects analysis; Groom et al. 2011b) developed the shade and temperature models that were generally used by the PSTM described here.

1.2 RipStream background

1.2.1 Study overview

RipStream was conducted at 33 sites in the Oregon Coast Range. Sites were situated along first- to third-order streams on privately owned (18 sites) and state forest (15 sites) lands dominated by Douglas fir (*Pseudotsuga menzeisii*) and red alder (*Alnus rubra*). Forest stands were 50-70 years old and were fire- or harvest-regenerated. Openings were dominated by shrubs such as vine maple (*Acer circinatum*), stink currant (*Ribes bracteosum*), salmonberry (*Rubris spectabilis*), and devil's club (*Oplopanax horridus*).

Criteria to select sites included an ability to collect at least two years of pre-treatment and five years of post-treatment data at every site, the inclusion of sites harvested with Riparian Management Areas (RMAs) that meet current FPA and state forest standards, minimum treatment reach lengths of 1000 feet, and assurance that the upstream "control" reach would remain unharvested for the duration of the study. Streams needed to qualify under the FPA as "Small" or "Medium" (mean annual streamflow < 2 cfs or between 2 and 10 cfs, respectively), and streams needed to be free of recent impacts from debris torrents and active beaver ponds. We obtained sites by requesting that industrial private and state forest managers in the Oregon Coast Range provide ODF with a list of stream reaches that met the criteria and would be harvested no sooner than 2004. An initial list of 130 stream sites was reduced to 36 sites that met study design constraints. Three more were subsequently dropped due to changes in harvest plans. While there was an initial attempt to exclude sites with beaver activity, a beaver dam ponded 722 ft of the 3,806 ft treatment reach for site 7801 during the first and second post-harvest years.

Each site has a control reach immediately upstream of a treatment reach. The control reaches were continuously forested to a perpendicular slope distance of at least 200 feet from the average annual high water level. Reach lengths varied from 450 ft to 6000 ft with means of 905 ft and 2244 ft for the control and treatment reaches, respectively (Dent et al. 2008).

1.2.2 Treatments

Forest Practices Act On Private Sites: Eighteen sites were established on private forest streams. Sites were harvested following contemporary FPA standards which require riparian buffers along fish-bearing streams to protect stream temperature, provide future large wood for streams, and retain other ecological services (Oregon Department of Forestry 2007). Measured as a slope distance, the RMAs are 50 and 70 ft wide around small and medium fish-bearing streams, respectively. Both small and medium streams have a 20-ft no-cut zone immediately adjacent to the stream. Harvesting is allowed in the remaining RMA to a minimum basal area of 40 (small streams) and 120 (medium streams) $\text{ft}^2/1000 \text{ ft}$. See section 2.6 for more details.

Oregon State Northwest Forest Management Plan (NWFMP) on State Sites: In Oregon, lands administered directly by Oregon Department of Forestry (state forests) are managed under FMPs for multiple resource objectives (e.g., recreation, wildlife) in addition to timber production and require riparian protections that exceed FPA minimum values. Fifteen sites were established on state forest lands. All but two RipStream state forest sites had RMAs managed according to the Northwest Forest Management Plan; two sites were managed according to the Elliott Forest Management Plan which has identical riparian management strategies. We therefore refer to the management of all 15 state sites as NWFMP. Measured as a horizontal distance from the stream's edge, state RMAs are 170 ft wide for all fish-bearing streams with a 25-ft no cut zone. Limited harvest is allowed between 25 ft and 100 ft of the stream only to create mature forest conditions. Further specifications are provided in section 2.7.

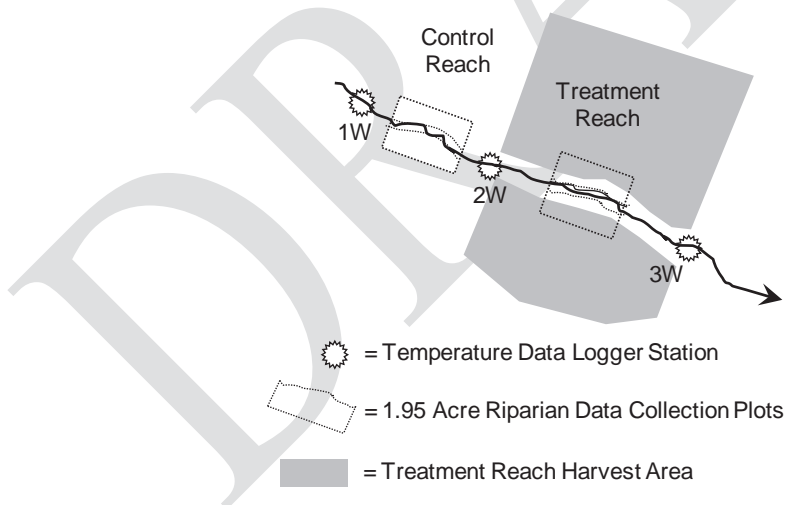


Figure 1. RipStream site layout. Control and treatment reaches are defined by the position of Station 2 and 3 temperature data logger stations that coincide with harvest boundaries. Two riparian data collection plots are situated midway along both the control and treatment reaches.

1.2.3. Data Collection

Optic Stowaway Temp and HOBO Water Temp Pro data loggers (Onset Computer Corporation, Bourne, Massachusetts) were annually deployed at four stations beginning in 2002 or 2003 (Figure 1). In the PSTM analysis we consider only Stations 1-3, not 4. 1W is at the upstream end of the control reach, 2W is located at the downstream end of the control reach and the upstream end of the treatment reach (i.e. shared logger) and 3W is at the downstream end of the treatment reach. Temperature loggers were deployed in shaded locations where stream flow was relatively constant, with reliable summer depth, and a well-mixed water column. Logger accuracy was checked prior to installation, in the field, and following retrieval with National Institute for Standards and Technology-calibrated digital thermometers according to the stream temperature protocol described by the Oregon Watershed Enhancement Board (1999). Although both types of loggers we used are listed as ± 0.2 °C accuracy, we found that for over 500 pre- and post- deployment assessments in only two instances did loggers register errors of >0.1 °C. Daily temperatures that exhibited increases in diel fluctuation and increases and decreases in daily maximum and minimum temperatures that were not reflected in other probes during the same year or at the same location during other years were interpreted as influenced by air temperature and excluded from the analysis.

Stream channel data were collected at 200 ft intervals within each reach. Data included wetted width, bankfull width, thalweg depth, and stream gradient collected according to the protocol described by Kauffmann and Robison (1998). Stream shade was quantified at these intervals using a self-leveling fisheye lens digital camera (Valverde and Silvertown, 1997). Shade values were measured once pre-harvest and once post-harvest. Fish-eye photographs were taken in the middle of the stream, 1 m above the water level, and oriented due north. Shade values were calculated from the photographs using HemiView™ 2.1 software (Delta-T Devices, Cambridge, UK) as one minus the June 30 Global Site Factor (1 -GSF). The GSF is the proportion of both direct and diffuse energy under a plant canopy relative to the available direct and diffuse energy for the given site's latitude/longitude. Shade and gradient values were averaged for each reach.

Vegetation data were collected in four 500 by 170 ft plots on both sides of a study stream in the treatment and control reach (Figure 1). Plots were centered midway along each reach and contoured according to stream curvature. Vegetation plot data describe understory, overstory, downed wood, blowdown, and snag characteristics. The original purpose of including extensive vegetation plot data collection was to assess large wood recruitment, shade, and riparian structure following timber harvest. For this analysis we use a portion of the available riparian structure data (e.g., blow down, tree heights, basal area, species). Within each plot all living trees with a diameter at breast height (DBH) > 6 inches were tallied by species; each tree's distance to the stream was additionally recorded. Height, live crown ratio, and crown class (e.g., dominant, co-dominant, intermediate, overtopped) were additionally measured for 20% of the trees. Figure 2 depicts the plot layout that determined data collection. Tree data were recorded in 100 x 170 ft "lines". In the middle of each line was a transect, along which riparian downed wood was tallied. Each of the 1/100th acre subplots (circles) within each line were situated at 25 ft increments beginning at the stream's edge. Within each of a plot's 30 subplots,

contractors recorded landform, hillslope (measured towards the stream), shrub composition and percent cover, and seedling data.

Data were collected in all four vegetation plots per site pre-harvest and re-measured in harvested treatment plot or plots (if one or both stream sides were harvested, respectively) post-harvest. Blowdown was quantified in all plots post-harvest.




Figure 2. Data collection schematic for a single vegetation plot. Plots contoured stream channels. Within each plot are five 100 ft x 170 ft lines. Each line has a transect, and along each transect at 25 ft intervals are 1/100th acre sub-plots.

The PCW, effects, and PSTM analyses are limited to all pre-harvest data and data from the first and second post-harvest years. Data collection included hourly stream temperatures (collected annually between July 1 and September 15) and both channel data (for a complete list see Dent et al., 2008) and riparian vegetation data (overstory and understory) during the first years pre-harvest and post-harvest.

1.3 PCW analysis

Our publication on the PCW analysis (Groom et al. 2011a) presents an unusually complex analysis. The complexity stems from our goal of answering a regulatory question using a statistically valid analysis approach constrained by the structure of the regulation. Specifically, we wanted to know if stream temperatures observed in RipStream sites indicated whether timber harvest was increasing stream temperatures by more than 0.3 °C.

The Oregon Department of Environmental Quality water quality rules include several stream temperature criteria. Among them is the Protecting Cold Water (PCW) criterion that represents a federally-required antidegradation water quality rule component. The PCW applies to “cold” streams with temperatures below specific temperature thresholds

(DEQ, 2004). Anthropogenic activities such as timber harvest are not permitted to increase stream temperature by more than 0.3 °C above its ambient temperature.

The criterion is fairly straightforward to assess at a point source, such as below the outlet of a pipe. Samples would be taken immediately above the outlet and below the mixing zone past the outlet. However, forest harvests affect stream temperatures as a non-point source process. The PCW temperature metric of interest is the 7DMM, or seven-day mean of daily maximum temperatures. This value is quantified as a moving mean across a season. That is, the first 7DMM represents day 1 through 7, while the next 7DMM will represent days 2 through 8. The Protecting Cold Water criterion indicates that if any 7DMM is elevated by more than 0.3 °C above a baseline condition then the criterion is not met.

We performed an analysis of the PCW (Groom et al. 2011), reducing our hourly data to daily summaries, and then to 7DMM temperatures. To determine if any single 7DMM is out of compliance required us to establish a reference condition and then ask if any 7DMM during the test condition exceeded the expected amount by more than its prediction interval plus 0.3 C. For the detailed methods of the analysis, please see Groom et al. (2011a). We restricted the parameterization of the PCW analysis to adhere to rule language. The rules do not consider factors such as the length of the harvested reach, year effects, gradient, or other pertinent variables. Therefore our analysis similarly excluded these potential factors as well. In addition, the analysis examined every summer 7DMM datum relative to expected values. As a result, the analysis provided an answer to our analysis question (we saw an elevation of temperatures in private sites above the 0.3 °C threshold), but the analysis itself was convoluted enough that using it to obtain temperature change magnitudes would be meaningless or at best provide suspect results.

The results of the PCW analysis were presented to the Oregon Board of Forestry several times between 2009 and 2011. At the January 2012 meeting the Board of Forestry determined that the RipStream findings indicated forest practices contributed to the degradation of a natural resource, cold water. This finding triggered an FPA riparian rules analysis that potentially leads to a change in the riparian rules. This decision led ultimately to the development of the analysis approach presented here. However, the modeling effort and results of the PCW analysis are not incorporated in the PSTM due to factors discussed above. Instead, the PCW findings led to a second analysis, the effects analysis, which examined the magnitude of temperature increase at sites as well as important site variables that were related to observed temperature change. The effects analysis serves as the basis for the PSTM.

1.4 Effects analysis

In 2011 we published a second manuscript (Groom et al. 2011b) that delved into site variables and how they related to observed temperature change. This “effects analysis” abandoned the constraints of the PCW analysis. The effects analysis examines the contributions of different variables at explaining observed temperature changes, including

treatment reach length, average treatment reach shade, elevation, average treatment reach gradient, east-west deviation in degrees of the treatment reach valley (a simplified version of valley azimuth), change in control reach temperature, and watershed area calculated at 3W. We also included state and private ownership and the harvest status or whether a temperature measurement occurred during a pre-harvest or post-harvest year.

Combinations of these variables were used in a suite of mixed linear regression models.

We determined that mixed models were called for, as such models allow site intercept and the parameter value associated with the control reach temperature change to differ by site, accommodating random site-specific characteristics. Using the same mixed parameter in the models we compared the performance of 18 *a priori* temperature models (Groom et al. 2011b Table 1).

We examined model performance for four temperature change metrics. We summarized hourly stream temperature data to provide daily maximum, mean, minimum, and fluctuation (maximum-minimum) values for each station. We were interested in detecting changes in stream temperature due to site factors including harvest. We therefore defined the response variable as the daily difference between treatment reach 2W and 3W. To reduce analysis complexity we computed the average of this difference over a forty day period for each year (July 23 to August 15). This represents the time frame when we had the greatest number of functional loggers recording temperatures during a central portion of the summer months when maximum temperatures are observed in the Oregon Coast Range. We compared the 18 *a priori* temperature models against these four temperature metrics.

The PCW analysis described in section 1.3 focused on maximum daily temperatures, averaged across seven day periods. Although the effects analysis examined a suite of temperature metrics, our current effort is focused on predicting effects of harvest on daily maximum temperatures as the DEQ temperature criteria focus on these quantities.

Fortunately, the effects analysis' metric for maximum daily temperatures averaged over a 40-day period (40-day max) is virtually identical to the mean of the 7DMM values taken for the same time period (Appendix 1). Therefore, we interpret the 40-day max values as a suitable substitute for an average response of 7DMM values.

The analysis focuses on an average response of the 7DMM instead of individual values out of a need for model simplicity. Streams temperatures differ in how they change temporally and spatially. Had we used 7DMM values we would have needed to include some modeling component of 7DMM values to account for site- or even year-specific trends. The model would have also needed to incorporate autocorrelation and moving average corrections for these data. This level of complexity was avoided by using the average. A consequence, however, is that the 40-day max values do not capture the highest (or lowest) 7DMM values at sites.

The 40-day max values for all years pre-harvest and the first two years post-harvest at all sites were best explained by variables for the change in control-reach temperature (Control), treatment reach length (TRLengh), the mean of treatment reach shade measurements (Shade), and the first quartile of stream gradient measurements (Grad1Q).

Control values were the 40-day max temperature change for a given year between the 1W and 2W probes. The values were normalized by control reach length. TRLength was determined from field measurements of the distance between 2W and 3W. Shade was measured as described in section 1.2.3. Gradient values appeared generally skewed by high-gradient values, so we conducted the analysis on the first quartile (lower readings) of channel gradient, relating to slower stream flow rates.

The formulation for the best-supported statistical model was:

$$[1] \Delta T_{3-2ijk} = \alpha_0 + \alpha_j + (\beta_1 \Delta T_{Control_{2-1i}} + \beta_j \Delta T_{Control_{2-1i}}) + \beta_2 TRLength_j + \beta_3 Shade_k + \beta_4 Grad1Q_j$$

Subscripts indicate year data i , site j , and pre- or post-harvest status k . The model includes mixed-effects parameters for the linear model's intercept (α_j) and the slope value for Control temperatures (β_j). These values allow a different intercept and Control value for each site, assisting in accounting for the lack of independence between 40-day max observations at a single site. The modeling effort determined that providing the structure of a mixed model was advantageous over an intercept-only mixed model or a standard linear regression.

Pre-harvest shade was constant enough between values of 80% and 95% that few variables could account for between-site differences. We therefore created linear models to describe only post-harvest shade. We constructed the models as weighted linear regression due to the different number of shade measurements taken per treatment reach. Some sites had five or six measurements, others had over twenty. Therefore the variability in the models should account for variability in the number of shade measurements. We performed a logit transformation of shade values to address to cope with the limited dependence variable (shade values cannot exceed 1.0) and as a means to linearize data that exhibited a curvature (see Section 1.6.2). Figure 3 illustrates the effects of such a transformation.

In the shade model-selection procedure, we considered models that included valley azimuth. These models did not perform well; the effects of valley azimuth on shade should have already been taken into account by our hemispherical photographs.




Figure 3. A display of the effects of performing a logit transform of the shade variable (shade values in this study fall between 0.5 and 0.95).

For logit-transformed shade data the best-supported statistical model was:

$$\begin{aligned}
 [2] \quad \text{LogitShade}_{\text{Post}} &= \alpha_{\text{Shade}} + \beta_{1\text{Shade}} \text{BasalArea}_{\text{Post}} + \beta_{2\text{Shade}} \text{TreeHeight}_{\text{Preharvest}} \\
 &+ \beta_{3\text{Shade}} \text{BasalArea}_{\text{Post}} * \text{TreeHeight}_{\text{Preharvest}}
 \end{aligned}$$

The logit of shade post-harvest depended upon the basal area post-harvest at each site, tree height values measured pre-harvest, and their interaction. The explanatory variables were assessed within 100 ft of the stream, a distance that we anticipated would include most if not all of the tree foliage that contributed to shading the reach during the middle of the day, when solar radiation is at its maximum. Basal area was calculated as the mean basal area (m²/ha) of the two plots out to a horizontal distance of 100 ft, counting all living trees with diameters at breast height of 6" or greater. Mean pre-harvest height within 100 ft was 84 feet.

Main findings for the temperature modeling procedure included finding that shade (which, according to [1], changed values pre-harvest to post-harvest) is a major contributor in the model towards explaining model variation, with less shade associated with greater temperature increases between 2W and 3W (Table 1). The fixed value for Control was negatively associated with changes in the treatment reach temperature. The greater the change in stream temperatures between 1W and 2W, the greater an opposite change would be seen in the treatment reach. That is, if there were substantial cooling in the control reach in a given year relative to other years, the model would predict that even in the absence of harvest there would be a relative increase in the temperatures for the treatment reach. The power of the fixed and mixed-effects aspects of Control at influencing model fit is demonstrated in Figure 4. In particular, note that the cross-hairs (predicted temperature values) often are close to or essentially on top of the circles (observed values) even when the values are fairly far apart. This is an indication of the degree by which the model adjusts for year-to-year variation in stream temperature behavior.

Table 1 (From Table 4, Groom et al. 2011b). Fixed- and random-effect parameter values for a linear mixed-effects model and its associated temperature response variables. Parameters for treatment length and gradient are expressed as change in temperature per 1 km of distance or elevation. Observations = 119, Groups (Sites) = 33.

Fixed ^a	DF	Value	SE	p
Intercept	29.1	0.494	0.125	0.001
Control	21.5	-1.232	0.459	0.014
TRLength	28.2	0.800	0.304	0.014
Shade	94.5	-5.866	0.572	0.000
Grad1Q	30.3	-0.076	0.036	0.040

Random	Std.Dev
Intercept	0.441
ControlTemp	3.564
Residual	0.079

^aControl reach temperature change = CT, gradient = GR, shade = SH, treatment length = TL.

^bFor Diel Fluctuation the variable for GR is replaced by elevation (EL). Other parameters in the model are the same.

Stream gradient (Grad1Q) and TRLength are associated with a change in temperature, with shallower gradients and longer treatment reaches associated with temperature increases over the course of the treatment reach. We interpret these findings to relate to the amount of time water within a stream is exposed to increased solar radiation in the harvest reach. Overall, when examining observed or partial residual temperature values for pre- and post-harvest periods, we found no temperature increase for State forest sites and ~ 0.7 °C increase for all private sites. A similar amount of temperature increase was predicted for a temperature model that used an indicator variable for ownership instead of a shade variable.




Figure 4, from Figure 3 in Groom et al. (2011b). Observed and predicted temperature changes for maximum temperatures (°C) by site. Pre-harvest and post-harvest observations are represented by open and filled circles, respectively. Each point represents one year of data collection at a site. The crosses represent predicted values from model Grad_Shade. Above each site's data is listed its site number, ownership ([S]tate or [P]rivate), post-harvest shade value, and in parentheses the change in shade value pre-harvest to post-harvest. Sites are ordered from the upper left to lower right by the observed change in shade values. Vertical differences of points within a pre-harvest or post-harvest category indicate a between-year change in the temperature relationship between 2W and 3W.

The shade model indicated that greater amounts of basal area post-harvest and shorter trees were related to greater amounts of shade. The model explained 0.69% ($R^2 = 0.69$) of

the variation ($n = 33$, $\alpha_{Shade} = 1.795$, $\beta_{1Shade} = 3.100e^{-2}$, $\beta_{2Shade} = 6.250e^{-2}$, $\beta_{3Shade} = 4.680e^{-4}$, model $p < 0.001$). Other similar models were also explored.

1.5 Need for a more flexible approach – Bayesian modeling

The PCW analysis indicated that temperature increases above threshold levels were occurring on privately-owned timberland, and the effects analysis provided two models that linked temperature change within a reach to shade (eqn. 1), and shade values post-harvest to basal area and tree height (eqn. 2). We saw an opportunity to use the models from the effects analysis to predict the amount of basal area necessary to maintain shade levels, which in turn would be expected to prevent harvest-related stream temperature increases.

We needed a means to run the temperature and shade models backwards to reach predictions. That is, we wanted to provide basal area values for the shade model, obtain an expected shade amount, and then enter this shade amount into the temperature model to obtain a temperature change prediction. A major hurdle in performing such a procedure was that both of the shade and temperature model relationships contained error terms. We did not see a clear way to effectively join these models and rigorously account for error.

We became aware of a modeling approach called Bayesian Hierarchical Modeling. The idea behind this variety of modeling is that different relationships that affect each other can be combined, with findings from one relationship informing the other. An example of this approach involves Antarctic fur seals (*Arctocephalus gazella*; Hiruki-Raring et al. 2012). The authors determine factors associated with fur seal pup mass based on different relationships between prey densities, sea ice, maternal mass, and several other values. The relationships among these variables ultimately assist in estimating effects on pup mass. From our effects analysis, we have a temperature and shade model. The models do share a commonality: at least for post-harvest periods, shade appears as a dependant variable in the shade model and an independent variable in the temperature model.

Bayesian modeling in general is different from the modeling approach (Frequentist) used in the PCW and effects analysis. A central difference is the philosophical and analytical approach taken towards data analysis. In a Frequentist analysis data are assumed to be random expressions of fixed parameters (Ellison 1996). That is, if the world were virtually static and we could exactly repeat the same experiment many times, the processes affecting the data would be fixed, but like a coin flip the data would be similar, but not identical, to previous realizations of the experiment.

From a Bayesian perspective, it is the parameters that are treated as random and the data as fixed. The data resulted exactly the way they were because of myriad forces in the world. In the fixed-world example, if the experiment were run again the data would be identical. The various shifting effects in the world in turn make it so that our crude parameters of interest appear to follow random distributions.

There are some distinct advantages to using a Bayesian analysis.

- 1) Models may be hierarchically combined. That is, since our shade and temperature model share information (the Shade variable), the models could be joined to estimate values simultaneously while accounting for error associated with parameter fit and nuisance parameters.
- 2) Statistical models used in Frequentist analyses may be used for Bayesian analyses as well; it is the manner in which parameters are estimated that differ, not the structure of the parameters in the model. We therefore can make use of the shade and temperature models as formulated from the effects analysis.
- 3) Missing data may be estimated. Bayesian analysis works backwards and forwards simultaneously. While the analysis estimates dependent parameter values it can use the parameter estimates to create an estimate for the missing value.
- 4) Bayesian results more intuitive to understand than Frequentist results. Frequentist statistics speak to the probability of observed results given the repetition of data collection infinite times. A Frequentist 95% confidence interval tells us that the true value of interest (e.g., the true mean of a population) should fall within the estimated confidence interval 95% of the time if the study were repeated many times. A Bayesian 95% credibility interval (analogous to a confidence interval) around a mean indicates a 95% probability that the mean lies within the interval (Ellison 1996).
- 5) Bayesian models depend on incorporating knowledge. A Bayesian model must be provided with parameter priors; i.e., information regarding the size of an effect, its variance, and the expected distribution. If the magnitude of the effect is set to zero the prior is uninformative; Frequentist and Bayesian analysis results for the same model will be very similar if uninformed priors are used.
- 6) Obtaining derived estimates (e.g., predicted results given specified conditions) from a Bayesian analysis is relatively straightforward.

Bayesian analyses have some shortcomings:

- 1) Priors can influence model results if the priors are erroneous. This applies to the prior mean, variance, and distribution. If the variance is too restrictive or the distribution incorrect, the analysis may be prevented from correctly estimating values.
- 2) Convergence must be achieved. Bayesian analysis has only recently (10-20 years) become generally feasible as computer algorithms (i.e., Markov Chain Monte Carlo techniques) are relied upon to determine the posterior distribution of estimates run for $10^3 - 10^6$ iterations. As the iterations run the parameter values improve (fit the data better). The results of the analysis are the various parameter point estimates taken from some number (e.g., 1000) of these iterations after the parameter values have stabilized. Each set of iterations for a parameter is called a chain. Modeling techniques allow us to assess the behavior of multiple chains. If one or more chains have not arrived at the same range of values then the model has not converged. We adjust the number of iterations to ensure that we include only values from chains that have reached convergence. At issue is that, particularly for complex models, the overall set of chains may have converged on a local value, not the “true” value.

- 3) If there are too many parameters in the model the model may overfit the data, or explain each datum well but have little ability to extrapolate well. Cross-validation techniques like leave-one-out can assist in assessing this condition.

We wished to use the Bayesian model to create model predictions. Specifically, we wanted to simulate harvests of different prescription types on our plot data and determine how the changes in the plot structure would affect shade values (via the shade model) and in turn how those changes would affect the change in temperature between 2W and 3W. We knew we could employ our best-performing temperature model from the effects analysis. However, for the Bayesian analysis, we decided to revisit the shade model, as the previously described model had some characteristics we wished to improve on (discussed in 1.6.2).

1.6 Developing sub-models

The temperature and shade models from the effects analysis represent different varieties of linear regression models. The temperature model included mixed effects parameters while the shade model is a weighted regression. We created null-prior Bayesian versions of each model type to ensure that resulting estimates matched the Frequentist estimates. Once we had the Bayesian parameterization decided upon, we could proceed with combining the models (Section 1.7).

1.6.1 Temperature

As a check of a correctly-functioning Bayesian model, we examined eqn 1 in R (library nlme, function lme, method = REML) as well as JAGS and WinBUGS. The R values differ slightly from Groom et al. (2011b), as the published values were run in SAS. The results of the probabilistic model are as follows:

	StdDev	Corr
(Intercept)	0.6683557	(Intr)
c_ControlTemp	1.8203856	0.306
Residual	0.2827235	

Fixed effects: Response ~ Control + TRLength + Shade + GradIQ

	Value	Std.Error	DF	t-value	p-value
(Intercept)	0.476914	0.1246656	84	3.825543	0.0003
Control	-1.241839	0.4373531	84	-2.839443	0.0057
TRLength	0.828420	0.3006459	30	2.755467	0.0099
Shade	-5.923156	0.5647864	84	-10.487427	0.0000
GradIQ	-0.083204	0.0423497	30	-1.964699	0.0588

Note that all independent variable values were centered for the analysis. The Bayesian results for the same model produced virtually identical findings for the fixed effects and similar estimates for the random effects (Table 2).

Table 2. Bayesian estimates for the effects analysis temperature model using uninformed priors. Percentages are available for constructing credibility intervals (95% CI = between 2.5% and 97.5%).

	mean	sd	2.50%	25%	50%	75%	97.50%
Random Effects							
intercept	0.71	0.11	0.52	0.63	0.7	0.77	0.96
Control	1.94	0.46	1.16	1.62	1.9	2.22	2.94
Residual	0.27	0.23	-0.2	0.12	0.28	0.45	0.68
	mean	sd	2.50%	25%	50%	75%	97.50%
Fixed Effects							
Intercept	0.48	0.13	0.21	0.39	0.48	0.57	0.75
Control	-1.24	0.48	-2.17	-1.55	-1.25	-0.93	-0.29
TRLength	0.83	0.32	0.23	0.61	0.82	1.04	1.51
Shade	-5.91	0.59	-7.07	-6.29	-5.9	-5.51	-4.74
Grad1Q	-0.08	0.05	-0.17	-0.11	-0.08	-0.05	0.02

The parameterization for the Bayesian model is provided in Appendix 2 (Section A2.1).

1.6.2 Shade

As mentioned in 1.5, we wished to re-visit the shade model. The effects analysis shade model was limited to examining post-harvest basal area out to 100 ft, not the full 170 ft measured. Therefore, our ideal shade model would:

- 1) Explain results well
- 2) Make intuitive sense
- 3) Include all vegetation plot data out to 170 ft horizontal distance from the stream
- 4) Include a measure of distance between stream and harvest at each site

The shade model from the effects analysis performed well and made sense (i.e., post-harvest shade was related to post-harvest basal area and tree height). However, it was limited to examining basal area only out to 100 ft horizontal distance. We originally limited it to 100 ft because the trees were on average 84 ft tall (less than 100 ft, greater than 75 ft); since solar radiation is most powerful between 10:00 AM and 2:00 PM in the summers, we anticipated that trees beyond this distance would not provide substantial shade to the stream. For this analysis we wished to more conclusively verify this assumption, and potentially include a measure of harvest distance to enhance the model's performance and inform the development of alternate harvest scenarios. Therefore we sought an analysis that incorporated all vegetation plot data recorded out to 170 ft and included a measure of harvest distance from the stream.

Incorporating a measure of distance from the stream to the harvest in the analysis proved difficult. We required a measurement that reflected actual harvest boundary distances. The vegetation plot data provide information on the distance of each tree from the stream as well as the vegetation plot line it was found in. To validate this method, we compared it against visually-determined harvest distances based on cumulative plots of basal area as a function of distance. By superimposing pre- and post-harvest cumulative basal area plot we observed where the two lines appeared to diverge. Selecting this divergence point was a subjective exercise, so we placed our assessments for every treatment plot in 25 ft categories (e.g., harvest between 26 and 50 ft horizontal distance; Figure 5). Assessments were corroborated by examining the difference between pre- and post-harvest values at 5 ft increments (not shown).

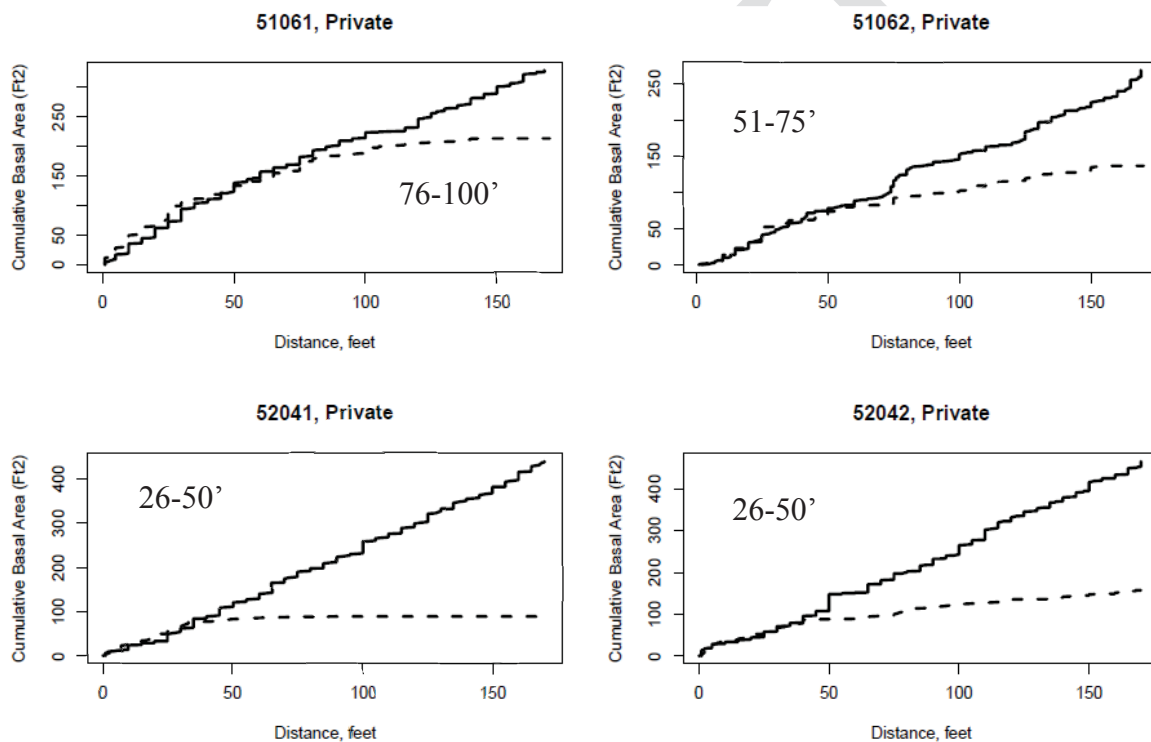


Figure 5. Cumulative basal area (ft²) of four individual treatment reach vegetation plots, as a function of distance. Solid lines represent pre-harvest data, dashed lines are post-harvest. Observed departure points between the two lines are between 76-100 ft.

Once we had the visual assessments of harvest distances we compared them to the mean of the distance to the outermost trees in each of a vegetation plot's five lines (MeanMaxDist). We examined the relationship between the two harvest distance assessments to verify that MeanMaxDist approximated the visually-determined harvest distance (Figure 6). We interpreted the R^2 value of 0.88 as indicative of an essentially good fit. Therefore, MeanMaxDist could serve as a measure of harvest distance for our analysis.

We examined suites of models including basal area and MeanMaxDist. An interaction

model including MeanMaxDist performed well:

$$\begin{aligned}
 [3] \text{Shade}_{post} &= \alpha + \beta_1 \text{BasalAreaPost}_{170} + \beta_2 \text{MeanMaxDist} \\
 &+ \beta_3 \text{BasalAreaPost}_{170} * \text{MeanMaxDist} \\
 &+ \beta_4 \text{TreeHeight}_{preharvest} 170
 \end{aligned}$$

Diagnostic plots indicated that it generally conformed to linear model assumptions, and had an R^2 of 0.71. The β_2 estimate was not significant but the interaction term (β_4) was. Although this model seemed promising, individual variables involved did not appear to exhibit a linear fit with shade (Figure 7).




Figure 6. Comparison of vegetation plot distance empirical measurement (MeanMaxDist) and a visual assessment of harvest distance based on cumulative basal area by distance (MeanDistBoth). The line represents the linear fit of the two variables. Each point is the mean value of the two treatment vegetation plots at a site.

Shade appears to have a non-linear relationship with basal area of plots out to 170 ft and MeanMaxDist (Figure 7A, B). The relationship appears tighter for basal area and shade. This non-linear relationship encouraged a deeper consideration of how we were modeling shade. In particular, Figure 7B appears to have two different slopes on two different intervals in the relationship between post-harvest basal area and shade, with slopes changing at $\sim 150 \text{ ft}^2$. Prior to that point the relationship is steep, afterwards less so. Our

original assumption that trees may influence stream shading out to a certain distance from a stream appeared supported. We therefore examined a different way to consider the MeanMaxDist data in the analysis.

We suspected that harvests that were distant from the stream would have little impact on shade while those that were closer would have greater impact. At the same time, the closer the cut to the stream, the less basal area post-harvest. We therefore examined the relationship in Figure 7B by asking which of the basal area points corresponded to cut distances of ≤ 75 ft, 100 ft, and 125 ft from the stream (Figure 8). We indicate which points fell below or above the “cutpoints” and fit a line to each set of points. The slope of the green line (above the MeanMaxDist cutpoints) becomes shallower from Figure 8A to Figure 8C. The slope of the orange line (below the cutpoints) is virtually identical for Figure 8A and B and then becomes shallower and fits the below-cutpoints worse in Figure 8C. We interpret these figures to indicate that basal area information is most relevant for sites harvested within 100 ft of streams, but that the information quality is degraded if we look out to 125 ft.

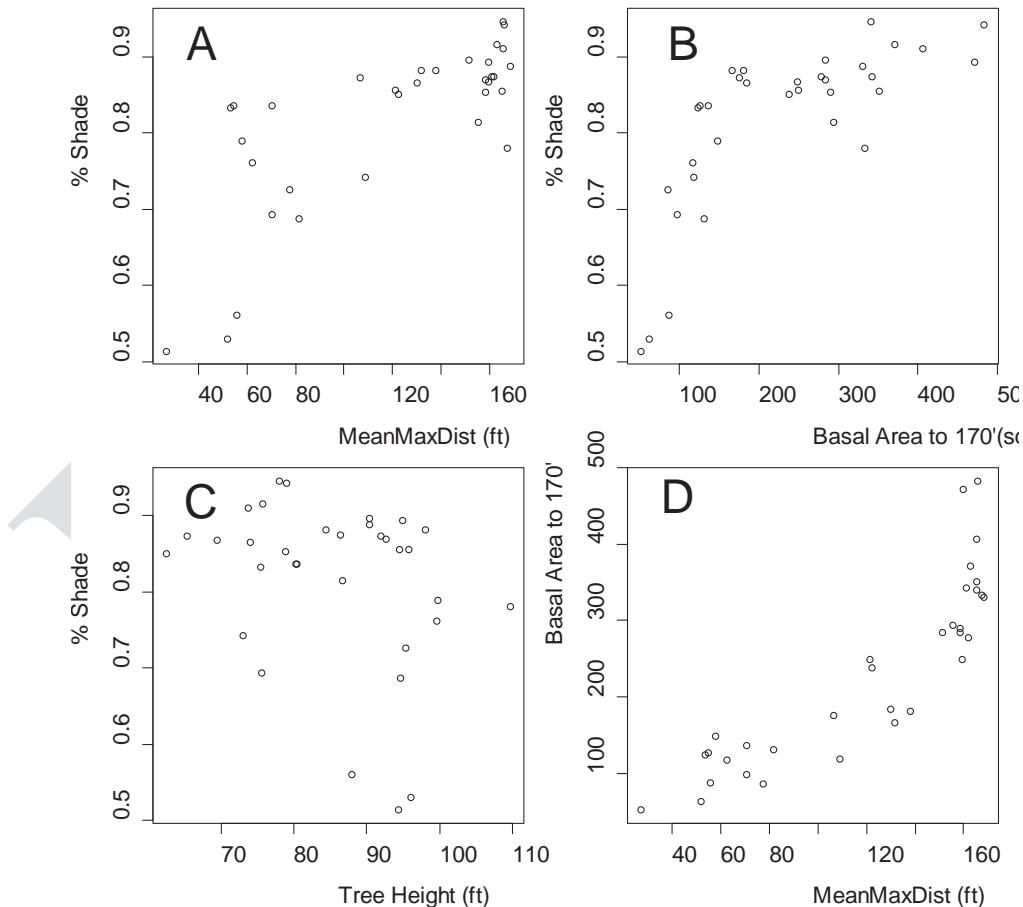


Figure 7. Scatterplots of percent shade post-harvest and MeanMaxDist (A), basal area of plots out to 170 ft (B), and pre-harvest tree height (C). The relationship between MeanMaxDist and basal area are plotted in (D).

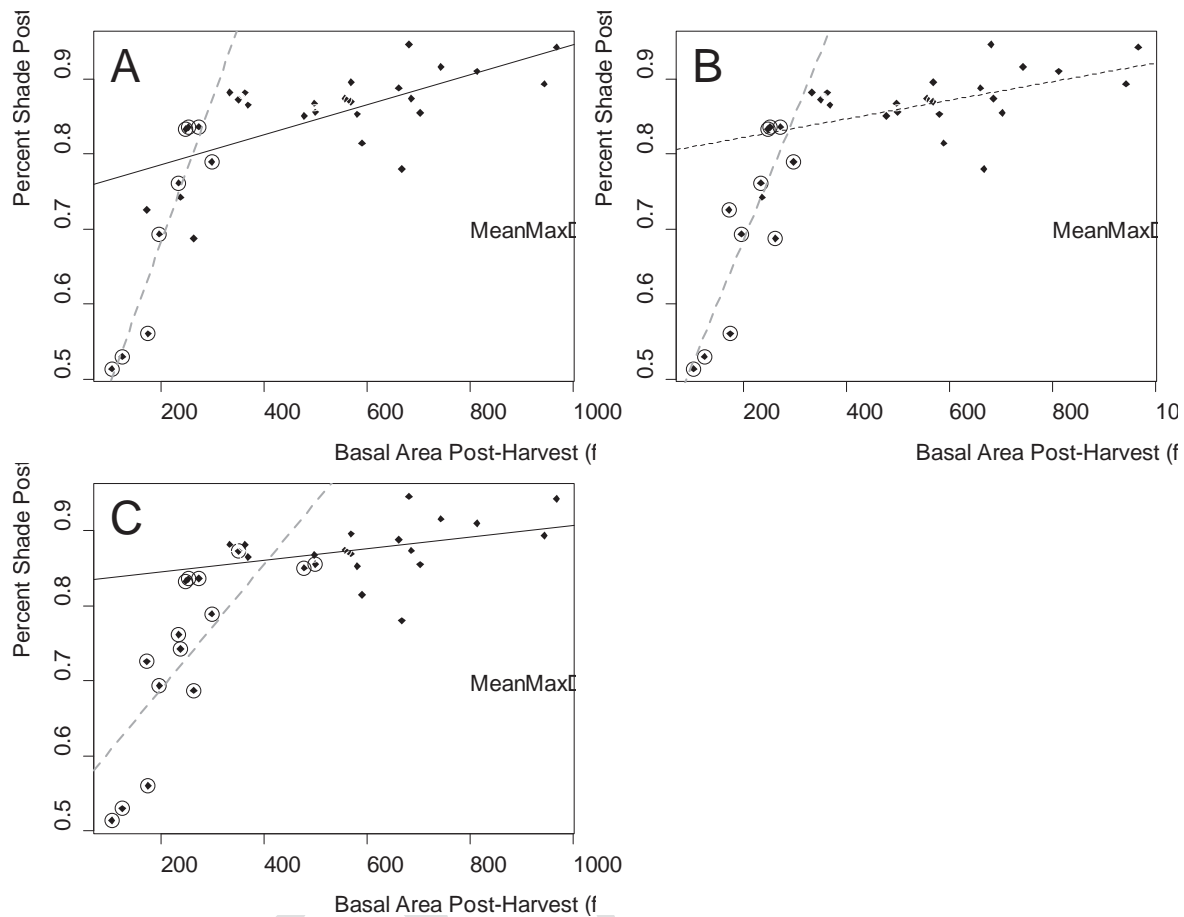


Figure 8. Fit of basal area post-harvest values ($\text{ft}^2/1000 \text{ ft}$) to stream shade while parsing data by the distance to harvest from the stream. Figure 8A has circles around all site data where the variable MeanMaxDist $\leq 75 \text{ ft}$. Figures 8B and 8C do the same for MeanMaxDist $\leq 100 \text{ ft}$ and 125 ft respectively. The solid black line is fit to points above the MeanMaxDist cutpoint, while the dashed grey line is fit to those data below the cutpoint.

We found further evidence that basal area beyond 100 ft was not useful for explaining shade. A line fit to the $< 100 \text{ ft}$ distance sites (Figure 8B) had a steep (and statistically significant) slope, indicating a relationship between shade and basal area post-harvest for sites with harvest distances $< 100 \text{ ft}$. In contrast, the slope of the shallower (green) line is not statistically different from zero. A separate analysis that compared a model that limited inclusion of post-harvest basal area to within 100 ft of a stream to a model that included basal area to 100 ft and from 100 ft to 170 ft (two basal area measurements in the same model) indicated that the more simple model was preferable and that no information was gained by including trees from 100 to 170 ft ($\Delta\text{AIC} < 2$). We therefore interpreted these findings as justifying the inclusion of trees no further than 100 ft horizontal distance from the stream in the shade analysis.

With this decision we focused on shade model selection. Previously, in the effects analysis, we examined a model with the logit of shade as the dependent variable. Was this

still supported? Comparing Figure 9 A and C to B and D (same independent variables, different dependent variables), it appears that there may be less of a curve evident in the data when shade is logit transformed and the independent variable is percent of basal area removed (Figure 9C). We therefore decided to use the logit of shade as the dependent variable. We used model selection to determine independent variables for inclusion. (Note: a quadratic term did not accommodate the curvature in 9B and 9D as well as the logit transformation.)

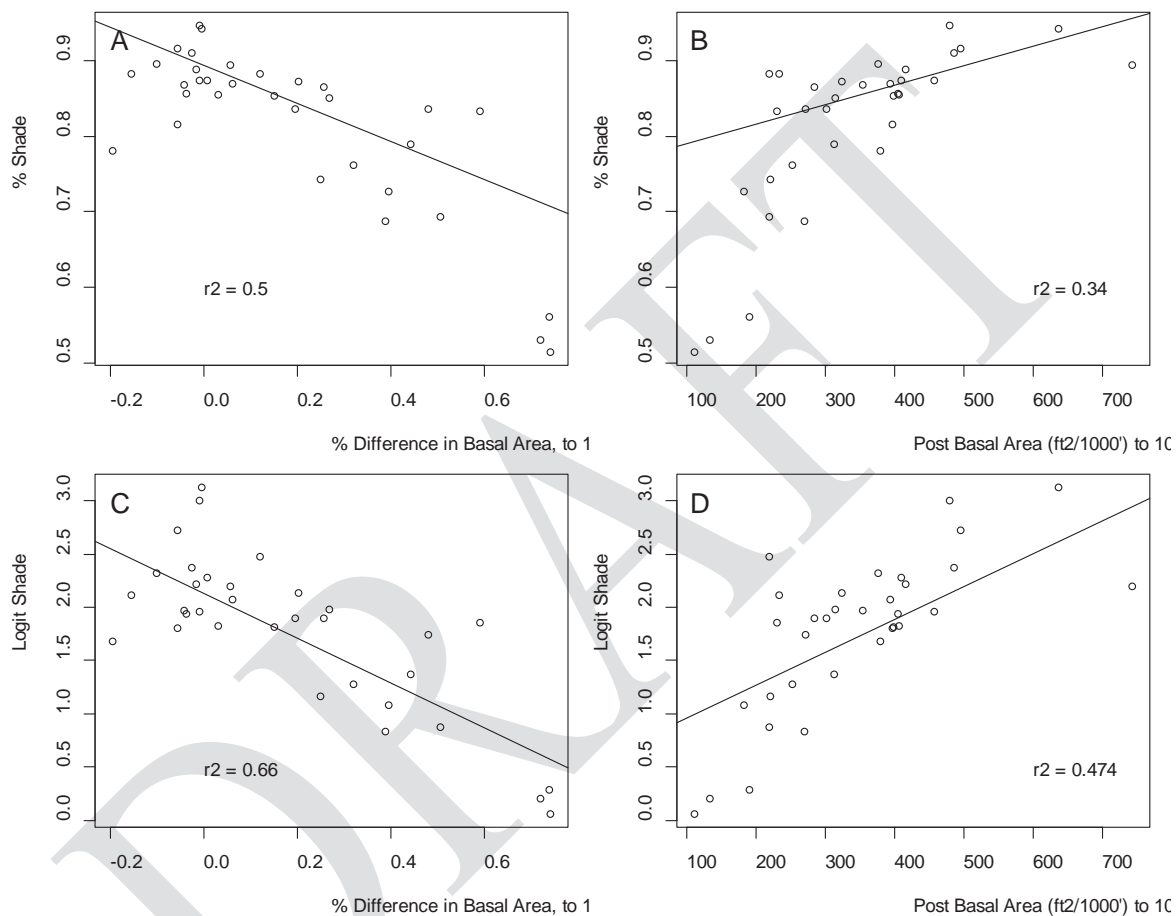


Figure 9. Plots comparing the fit of dependent variables shade (A, B) and the logit of shade (C, D) against independent variables “% difference in basal area, to 100 ft” (A,C) and “post-harvest basal area to 100 ft” (B,D). Lines represent weighted regression fits, R^2 values are provided.

We additionally wished to re-visit the metric for post-harvest basal area. Although post-harvest basal area originally performed relatively well as a predictor, it conveyed no information of pre-harvest stand information. Sites were generally well-shaded pre-harvest (figure 10). Yet, they differed in pre-harvest basal area values. We reasoned that a set reduction in basal area may disproportionately affect stream shading at sites with lower pre-harvest basal area values. Therefore we created the variable “Percent Basal Area Reduced,” which is, for all trees within 100’ horizontal distance, the amount of basal area

pre-harvest minus the basal area post-harvest divided by the pre-harvest value.




Figure 10. Appears as Figure 4 in Groom et al. 2011b. Plot of average treatment reach shade values (%) for each site, grouped by harvest status (pre-harvest, post-harvest). On the left are state forest shade values, on the right shade are private forest shade values.

For the independent variable model selection we investigated model fitting with combinations of the independent variables: post-harvest basal area to 100 feet, the percent basal area reduced, mean pre-harvest tree height of trees to 100 feet from the stream, and the percentage of hardwood pre-harvest within 100 feet, assessed by basal area. Models are presented in Table 3. The top six models included the variable “percent basal area reduced” which outperformed the variable for post-harvest basal area (present in the following six models). The top three models received similar AIC values (maximum $\Delta AIC = 2.53$) and the probability that one of the three models were the best of the set was 0.99. Given that the third-best performing model had an AIC value slightly greater than the variable penalty term for the top-ranked model (which had one extra variable) we essentially see the explanatory power of all three premodels as roughly equivalent. Because we seek a generally robust model (fewest parameters) we select model 3 as the preferred model.

Table 3. Model selection for fitting logit of shade values. k is the number of estimable parameters in the model, ΔAIC is the change in AIC between a model and the model with the lowest AIC value, and ω is the model weight (probability of being the best of the set). Models are sorted by ΔAIC . All independent variables were assessed to a distance of 100 ft from the stream. PctBA_red is the percent basal area reduced, PostBA is the post-harvest basal area, Ht is average tree height pre-harvest, and PctHWD is the percent of hardwood by basal area pre-harvest.

Model	Independent Variables	k	ΔAIC	r^2	ω
1	PctBA_red + PctHWD + PctBA_red * PctHWD + Ht	5	0.00	0.81	0.56
2	PctBA_red + Ht + PctBA_red * Ht + PctHWD	5	1.46	0.80	0.27
3	PctBA_red + Ht + PctHWD	4	2.53	0.78	0.16
4	PctBA_red + PctBA_red ²	3	10.37	0.70	0.00
5	PctBA_red + Ht	3	11.17	0.70	0.00
6	PctBA_red	2	13.16	0.66	0.00
7	PostBA + Ht	3	13.34	0.68	0.00
8	PostBA + Ht + PctHWD	4	15.07	0.68	0.00
9	PostBA + PctHWD + PctBA_red * PctHWD + Ht	5	16.03	0.69	0.00
10	PostBA + Ht + PctBA_red * Ht + PctHWD	5	16.10	0.69	0.00
11	PostBA + PostBA ²	3	18.04	0.63	0.00
12	PostBA	2	27.41	0.47	0.00

The equation for the selected model for site j , where all variables are considered out to a distance of 100 ft from the stream, is:

$$\begin{aligned}
 [4] \text{ LogitShadePost}_j &= \alpha + \beta_1 \text{PctBasalAreaReduced}_j + \beta_2 \text{PctHardwoodPre}_j \\
 &+ \beta_3 \text{TreeHeightPre}_j
 \end{aligned}$$

The observed vs. predicted fit for eqn 4 appeared essentially linear along the 1:1 line (Figure 11). Shade model diagnostics (Figure 12) indicate that residuals are fairly constant over modeled values, the data appear to be normally distributed, and that individual points are not exerting substantial leverage over the model. Points at the upper right of Figure 11 appear to be associated with greater standardized residuals relative to other points.




Figure 11. Weighted linear regression observed vs. predicted fit for the logit of shade (equation 4). The 1:1 line represents a perfect model fit; state forest sites are in red while privately-owned sites are in black.

The Frequentist fit for the estimates of the weighted shade model (Table 4) closely matches the Bayesian fit (Table 5). The measures of error differ, which appears to be an effect of the inclusion of sample weights. The parameterization for the Bayesian shade model is provided in Appendix 2 (Section A2.2).

Table 4. Output from R for modeling the logit of shade according to eqn. 4. The regression is weighted, with weights = $1/(\text{variance in logit shade})$.

	Estimate	Std. Error	t value	Pr(> t)
Intercept	-0.02507	0.05111	-0.491	0.62742
PctDiffBA	-2.30961	0.23724	-9.735	1.21e-10
TreeHt	-0.04406	0.01515	-2.909	0.00690
PctHWD	-0.74600	0.21445	-3.479	0.00161
Residual standard error: 0.5077 on 29 degrees of freedom				

Table 5. Bayesian estimates for the shade model described in eqn. 4 using uninformed priors. Percentages are available for constructing credibility intervals (95% CI = between 2.5% and 97.5%).

	mean	sd	2.50%	25%	50%	75%	97.50%
Intercept	-0.02532	0.10027	-0.22416	-0.09287	-0.0254	0.044207	0.166609
PctDiffBA	-2.31089	0.468524	-3.22638	-2.62181	-2.31706	-1.99762	-1.38907
TreeHt	-0.04382	0.029668	-0.10029	-0.06417	-0.04407	-0.02384	0.014772
PctHWD	-0.73969	0.421963	-1.55758	-1.0295	-0.74402	-0.4573	0.108633

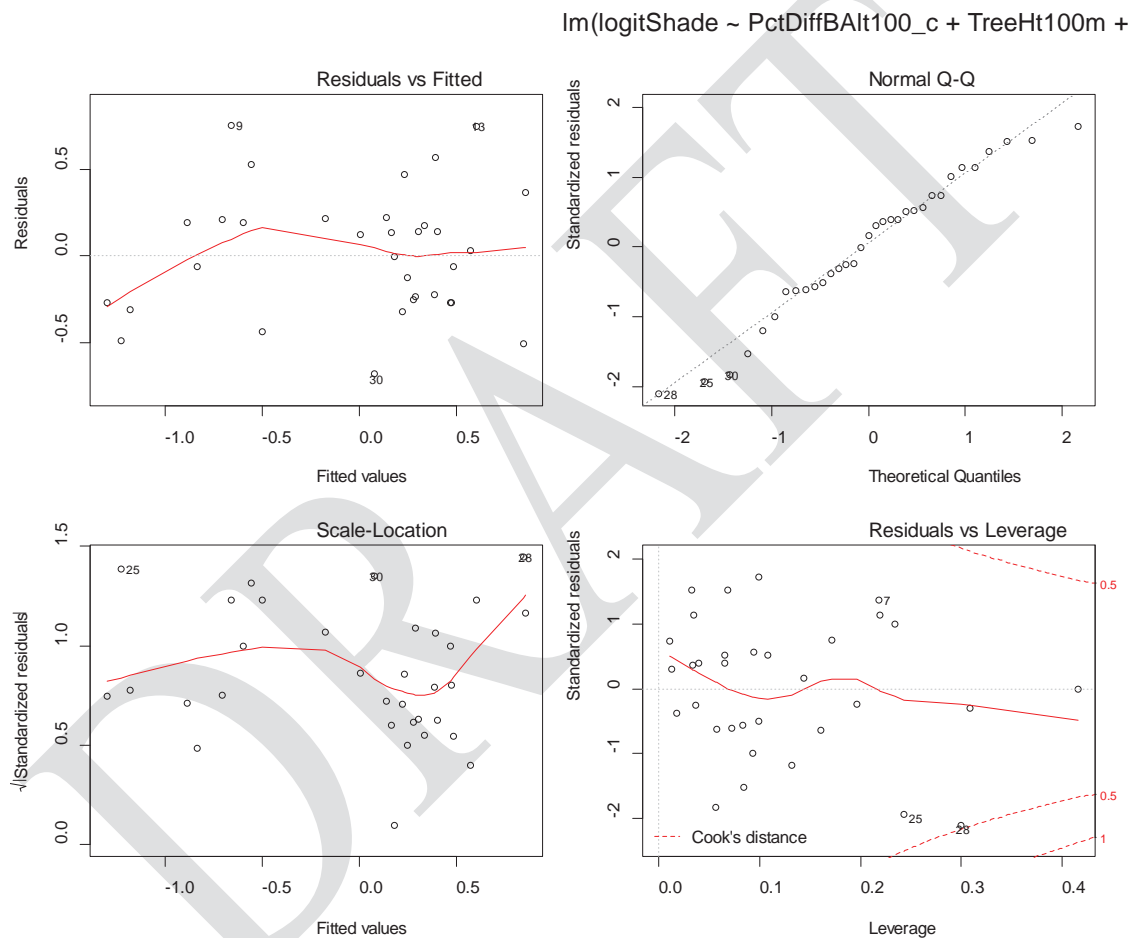


Figure 13. Diagnostics plots for the selected logit shade regression model (eqn 4).

1.7 Combining the sub-models

Linking the shade and temperature models was relatively straightforward. Our goal was to estimate post-harvest responses of shade and stream temperature to specified harvest prescriptions. Therefore, we did not need to alter the temperature model or its inputs for the pre-harvest period. Observed shade values were used as shade variable values pre-harvest. For the post-harvest period shade variable values were replaced with the

estimated post-harvest shade values (eqn. 4). This replacement was done to link the temperature sub-model and the post-harvest shade sub-model. The post-harvest shade sub-model itself was informed by post-harvest shade data. Both the observed and estimated logit shade values were transformed to percent shade (Appendix 2, Section A2.3).

[5]

$$\Delta T_{ij,k=1} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) + \beta_2 TRLength_j + \beta_3 Shade_{j,k=1} + \beta_4 Grad1Q_j$$

$$\Delta T_{ijk=2} = \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) + \beta_2 TRLength_j + \beta_3 Shade_{Post} + \beta_4 Grad1Q_j$$

$$LogitShadePost_j$$

$$= \alpha_{Shade} + \beta_{1Shade} PctBasalAreaReduced_j + \beta_{2Shade} PctHardwoodPre_j + \beta_{3Shade} TreeHeightPre_j$$

In Eqn. 5 ΔT is the change in stream temperature between 2W and 3W ($3W - 2W$); i = year, j = site, and k = timing relative to harvest (pre- or post-harvest). For clarification, where $k = 1$, $i = 1, 2, n$ years pre-harvest, where $k = 2$, $i = 1, 2$ years post-harvest.

Eqn. 5 contains all of the non-nuisance parameters estimated by the combined Bayesian model. Of note, parameter values in the shade sub-model have a *Shade* suffix, and we perform a logistic transformation of logit shade values and estimates for use in the temperature estimation equations. The temperature change sub-model maintains the same parameters pre- and post-harvest ($k = 1, 2$); therefore, their parameter distributions are estimated using data from both the pre- and post-harvest periods. Shade sub-model estimates are informed by post-harvest shade data and the temperature sub-model estimates by observed changes in temperature. This involves estimations of model error terms and other nuisance parameters (Appendix 2, Section A2.3). We use these estimated parameter distributions for the prediction scenarios.

Non-nuisance parameter estimates for [5] are presented in Table 6. Parameter estimates for the full model are similar but not identical to those presented in Tables 2 and 5. Of note, the credibility intervals for shade values have become narrower. For a more complete listing of parameter estimates, see Appendix 2, Section A2.4.

Table 6. Parameter estimates (abbreviated) for estimates of [5]. Sub-models are underlined. Sub-model means, standard deviations, and quantiles are presented.

Sub-models & Parameters	mean	sd	2.50%	25%	50%	75%	97.50%
<u>Temperature</u>							
Random Effects							
Control	1.921	0.440	1.186	1.599	1.887	2.199	2.861
Intercept	0.730	0.111	0.538	0.654	0.720	0.794	0.981
Residual	0.153	0.228	-0.297	-0.007	0.165	0.319	0.555
Fixed Effects							
Intercept	0.396	0.133	0.104	0.310	0.400	0.484	0.645
Control	-1.092	0.460	-1.951	-1.399	-1.105	-0.807	-0.152
TRLengh	0.871	0.336	0.206	0.663	0.878	1.077	1.511
Shade	-5.606	0.844	-7.341	-6.153	-5.590	-5.030	-4.046
Grad1Q	-0.077	0.049	-0.179	-0.109	-0.076	-0.044	0.014
<u>Shade</u>							
Intercept	-0.279	0.066	-0.407	-0.321	-0.279	-0.237	-0.148
PctDiffBA	-2.776	0.305	-3.428	-2.973	-2.770	-2.558	-2.223
PctHwd	-0.585	0.249	-1.092	-0.754	-0.583	-0.414	-0.100
TreeHt	-0.065	0.017	-0.100	-0.076	-0.066	-0.054	-0.031

We are interested in predicting the effects of specific harvest scenarios on the changes in stream temperature for treatment reaches. We do so by controlling as many variables as possible. Our approach is to predict a change in stream temperature using a harvest scenario resulting in a specific per-site value for Percent Basal Area Reduced (*PctBasalAreaReduced*).

$$\begin{aligned}
 [6] \quad \Delta \hat{T}_{i=1,j,k=2} &= \alpha_0 + \alpha_j + (\beta_1 \Delta TControl_{2-1ij} + \beta_j \Delta TControl_{2-1ij}) \\
 &+ \beta_2 TreatmentReachLength_j + \beta_3 (inverse \ logit \ of: \alpha_{Shade} \\
 &+ \beta_{1Shade} PctBasalAreaReduced_j + \beta_{2Shade} PctHardwoodPre_j \\
 &+ \beta_{3Shade} TreeHeightPre_j) + \beta_4 GradientQuartile_j
 \end{aligned}$$

Equation [6] is populated with the estimated parameters from [5]. Importantly, change in temperature for equation [6] represents a derived value, not an estimated value. The Bayesian model obtains all estimated values from [5]. For every scenario, we obtain two predicted temperatures from [6]. The first is the predicted change in treatment reach temperature for the first year post-harvest with a harvest effect (simulated or observed). The harvest effect is represented by the variable *PctBasalAreaReduced* and is calculated from vegetation plot data (see Section 1.2.3). The second prediction sets the variable *PctBasalAreaReduced* equal to zero change between pre- and post-harvest basal area. We subtract the second prediction from the first, with the difference representing the predicted

increase in temperature, per site, due to harvest. This procedure effectively controls for site-specific influence of non-shade temperature variables.

For a formulation of the model and incorporation of temperature prediction, see Appendix 2, Section A2.3.

1.8 Model evaluation

There were several aspects of model evaluation that we investigated. We wanted to determine model fit (i.e., how well did it predict observed values) and determine how well the model met assumptions. These evaluations do not consider simulated results because there are no observed values or “truth” against which to directly check predictions.

For model fit, we examined the first-year post-harvest estimated values as these were the values we wished to simulate. We plotted estimated vs. observed values (Figure 13). Of note, the estimated values represent the change in temperature between the probes 2W and 3W, taking into account all variables in the model including Percent Difference in Basal Area. We interpret the plot as indicating a linear fit with some under-estimation of temperature change at the larger increases in temperature.




Figure 13. Observed vs. predicted changes in treatment reach temperatures between probes 2W and 3W. The diagonal line represents a one-to-one fit; the blue diamonds represent the change in temperature between 2 and 3 W for the first year post-harvest.

We then examined model assumptions, some related to predicted values.

1.8.1 Assumption: Priors specification

We performed two evaluations of priors. The first was to examine the distributions of priors and determine if any appeared to have distributions truncated by prior specification. All prior distributions appeared to function well except for the uniform distribution for the nuisance parameter *idelS* (Appendix 2, section A2.2). As a uniform distribution, it estimated a value precisely at the distribution's midpoint (50 if from 0 to 100). We increased the possible range of values to 0-1000, and found that it predicted a mean of 500. Because this parameter was for the weighted shade regression model, we examined replacing the uniform distribution with a gamma distribution and checked model performance. We executed the shade analysis with the gamma distribution and null prior values for the mean and variance of 0.01 and 0.001. All four of these iterations produced very similar outcomes in other parameters (which in turn were approximately equivalent to the Frequentist shade model). We interpreted the original prior as functioning well.

The model contains many priors with Gaussian distributions. They all have starting means of 0, but typically have precisions (reciprocal of variances) of 0.001. We increased the variance to 0.01 for all parameters and ran the model for 2×10^6 iterations. Results were very similar between the two levels of precision. Most means of parameter distributions exhibited less than a 1% difference; the largest absolute difference was for the covariance parameter estimate (4%). We interpreted these results to indicate that the priors were appropriately specified.

1.8.2 Assumption: Parameter convergence is achieved

Models were run with 6 chains (i.e., 6 parallel estimation procedures) for a number of iterations beyond which we saw any steps in trace plots for any of the parameters. We cannot know if parameters settled for all six chains reliably at the same range of values but would have settled at different values for longer and longer runs. We were careful to include sufficient iterations to reach an estimation consensus.

1.8.3 Assumption: Model is not overfit

An overfit model is one which lacks generalization; it describes the data at hand well but would fail to predict useful responses given new inputs. One method for assessing this condition is to conduct a leave-one-out cross-validation. The idea is to see how well the model predicts a data point if that point were not included in the analysis. We conducted this analysis to re-create first-year post-harvest estimates in a similar fashion to Figure 13. Individual first-year post-harvest values were omitted. The Bayesian model imputed the missing values using the within- and among-site estimates and relationships.

An important consideration for the cross-validation approach is that we are conducting it for a mixed-effects analysis. When we leave out a data point, it affects the estimation of the random effects parameters for a particular site. This is especially true for sites where there may be a single value pre-harvest and post-harvest due to probe malfunctions. In these instances the mixed-effects parameter may resort to a global mean, altering the prediction greatly.

The leave-one-out cross-validation results are presented in Figure 14. The points are not as linear at higher temperature changes as in Figure 13. To an extent this should be expected, with less information available (particularly for random effects that are using at the most four or five data points to begin with). We have identified the data points from sites that were reduced to one point per site in the leave-one-out (LOO) procedure (red circles). The three biggest outliers are for an observed decline of 1.3 °C and observed increases of 2.1 °C. We interpret the LOO to indicate that the model sufficiently fits the data. Extreme outliers were generally those that had only one data point to use for random effects estimations, and if anything the model indicates an underprediction of temperature increase (conservative) relative to observed temperature increases.




Figure 14. Leave-one-out cross-validation of Bayesian model, first-year post-harvest data. Each point represents an estimated temperature response given the remainder of the temperature data set. Open triangles are state forest sites; closed diamonds are privately owned sites. Red circles surround data points that were estimated with only one data point informing random effects.

1.9 Model – specific assumptions

We performed assumption-checking procedures in Section 1.8 to verify the validity and performance of the model. Those procedures were standard for evaluating a predictive Bayesian model. However, there are some important assumptions for this specific model that affect interpretation and use of the model.

1.9.1 Inference and utility

Our model appears to fit the data well (Figure 13). However, its data requirements make it unlikely to suitably apply to other data sets. The model utilizes pre-harvest and post-harvest data for a treatment and control reach (with the control reach directly upstream of the treatment reach). It also depends on shade data and riparian vegetation information. It requires these data to feed a specific mixed-effects temperature sub-model and the associated shade sub-model. Therefore, we do not foresee the model being used directly outside of these study sites.

In addition, sites were not randomly selected. We used virtually all sites provided by industry and state forests that met our site inclusion criteria (including agreement to maintain an unharvested upstream control reach). Given our sample size (33 sites) and geographic extent of sites we interpret the results as representative for small and medium type-F streams in the Oregon Coast Range (three sites technically fell in the Interior georegion). Technically we lack the statistical inference a randomized study offers. On the other hand we are not aware of other randomized stream temperature studies that offer the same level of statistical power as this manipulative long-term study design.

1.9.2 Pre-harvest shade levels

Our model describes shade and temperature relationships for sites that all began with levels of shade that typically exceeded 80% (Figure 10; 4 sites had shade levels pre-harvest between 70 and 80%). It therefore may not describe well the thermal behavior of streams at sites subject to harvest that exhibit lower pre-harvest shade levels.

1.9.3 Harvest tree distributions

As described in Section 1.6.2 we excluded consideration of trees (basal area) measured farther than 100 ft from the stream, horizontal distance. However, within 100 ft of the stream we relied on percent basal area reduced, not the mean of the maximum distance to the edge of the harvest. Twenty-two sites had values for the mean of the maximum distance to the edge of harvest beyond 100 ft, so the variable was relatively non-informative when compared to percent difference in basal area.

Only one state forest site had substantial reduction in basal area within 100 ft of the stream; therefore, reductions in basal area predominantly reflected private ownership harvests. These harvests appeared as hard-edged clearcuts (i.e, not thinnings within the RMA). The variable for the percent difference in basal area within 100 ft of streams is therefore reflective of a hard-edged clearcut and not thinning. It is informed purely by changes in basal area within 100 ft from the stream; therefore, we assume in our simulations that harvest occurs from 100 ft from the stream and inwards which is reflective of our observations of site harvest patterns.

Part II: Harvest Simulation Approach

2.1 Vegetation data use

Vegetation plot data collection is described in Section 1.2.3. To summarize, we have 100% tree cruise information from all plots pre-harvest, from treated plots post-1, and blowdown were tallied in all plots during the post-1 period. We summarized the vegetation data from each plot to obtain different metrics (MeanMaxDist, basal area pre-harvest and post-harvest) for use in the shade analysis and to develop the metric Percent Basal Area Reduced.

To simulate a harvest we used pre-harvest plot data. These data were then subject to a specified harvest procedure (e.g., FPA harvest). We ran identical plot summary programs on the pre-harvest plot and the simulated harvest plot. One of the variables recorded for both scenarios (pre and post) was basal area within 100' horizontal distance of the stream. The horizontal distance is used in all cases as it is the metric the predictive analysis relies upon.

Plot vegetation data are further summarized into site data. Site vegetation characteristics are the mean of the corresponding treatment plot metrics. Therefore, if a simulated FPA harvest removed more trees from one bank than the other, the resulting value for basal area would fall between the two plot basal area values. We performed this step as the model itself was constructed with the vegetation data summarized in this fashion. As mentioned in Section 1.4, the model does not incorporate valley azimuth.

To obtain Percent Basal Area Reduced we subtract a site's post-harvest basal area (the mean of basal area values from the two treatment plots) from the site's pre-harvest basal area, divided by the pre-harvest basal area. This procedure was used to predict the effects of harvest under all scenarios except the as-harvested scenario. In the as-harvested scenario we used both pre-harvest and the measured (not simulated) post-harvest data.

Obtaining the measured post-harvest data for the as-harvested scenario required additional data manipulation. As mentioned above, only treated plots were fully re-measured post-harvest. If a site was harvested only on one side, we combined the treated plot with the pre-harvest data from the untreated plot. A further complication was blowdown. For estimating the effect on stream temperature, we removed plot blowdown from the basal area estimate. We omitted blowdown as it would likely reduce shade levels and we wished to predict as closely as possible the temperature increase due to the change in shade.

2.2 Predictions

2.2.1 As-harvested

As described in Section 1.7, we used the observed change in basal area between pre- and post-harvest to predict harvest-related changes in temperatures for individual sites (Figure 15). The privately-owned sites had predicted overall temperature increases of 0.93 °C while the state forest sites had predicted increases of -0.05 °C.

Within this prediction some of the State sites exhibited predicted temperature decreases. This was due to greater basal area within 100 ft of streams recorded post-harvest compared to pre-harvest vegetation cruise surveys. We attribute the increase in basal area due to tree growth, ingrowth of smaller trees, and measurement error.

2.2.2 State Forests

We developed an approach to simulate a state forest FMP riparian harvest of our pre-harvest treatment reach data (Appendix 3). All treatment reach pre-harvest vegetation data were reduced according to our interpretation of the FMP, simulating harvest on both banks (Figure 16). Sites that achieved mature forest condition, either by being dominated by hardwoods or by having many large conifers, were not reduced in basal area by the simulation (temperature increase = zero). Other sites received harvest within 100 ft as limited by minimum basal area and conifer number retention requirements. All other possible trees were removed. The analysis predicts that the average temperature increase for all sites subjected to a thorough FMP harvest is 0.19 °C, with a 95% probability (one-sided) that the mean would be less than 0.23 °C.




Figure 15. Predicted temperature increase due to change in basal area for all sites, as harvested. Boxplots indicate median values (mid-lines) within grey boxes (25% and 75% quartiles), with whiskers extending from the 2.5th – 97.5th percentile. Circles represent data outside of the whiskers. Sites are listed along with ownership (S or P for State or Private) and stream classification (Sm = small, Med = medium). The horizontal line indicates zero temperature increase.

2.2.3 FPA Harvest

We programmed a similar two-sided harvest of all sites, irrespective of ownership, to experience a complete removal of as many trees as permitted by the FPA (Appendix 4). This harvest relied on tree slope distance from streams. We then determined the percent change in basal area between the pre-harvest plots and the simulated FPA harvests of those plots and entered the value into the simulation portion of the Bayesian model. The results indicate a mean temperature increase of 1.77 °C with only a 5% chance that the temperature increase on average would be below 1.43 °C (Figure 17). This result indicates that 1) removal of all trees permitted under the FPA is predicted to result in significant warming, and 2) given the discrepancy between the as-harvested predicted warming and the simulated FPA harvest, a harvest as extreme as is portrayed in Figure 17 may not be common practice on industrial land ownership. We initially asked landowners participating in the RipStream study to extract timber from riparian areas as permitted by the FPA and NWFMP. This was done to ensure that the rules, not practices, were tested. We were concerned that a “business as usual” harvest would not allow us to test the rules (according to ODF’s 2002 compliance audit¹ 60% of landowners did not enter the RMA and on average operators left over 200% of the required basal area). Even with landowner

¹ ODF 2002. Best Management Practices Compliance Monitoring Project: Final Report. Oregon department of forestry forest practices monitoring program technical report #15.

cooperation, it appears that harvestable basal area may have remained.




Figure 16. Simulation of response of all sites to a two-sided State Forest FMP harvest. The “whiskers” (lines above and below boxes) identify the 95% credibility interval for each site (95% probability that the mean response falls between those lines). The boxplot on the far left, labeled “All,” represents the mean response of all 33 sites, and its 95% credibility interval is represented by the two dashed lines.

Riparian prescriptions in the FPA are to abide by rule language *averaged over 1000 feet*. Therefore, it is possible that the riparian plots, which are each 500’ long, captured at some sites metrics not representative of buffer dimensions or characteristics. Ancillary work examining orthophoto imagery and LiDAR on suspected aberrant sites provided no indication that vegetation plots were unrepresentative. Figure 18 displays the total basal area from pre-harvest data, as-harvested data, and the FPA and NWFMP simulated harvests. It appears that state forest and privately-owned sites typically harvested less than they potentially could have. All but one state forest site harvested well above our simulated NWFMP level. Approximately six out of the 18 private sites appeared to harvest at or almost at the FPA level. Twelve of the 15 State Forest sites appeared to essentially receive little or no entry within 170’ of the streams at the location of the vegetation plots.

The RipStream vegetation plots were cruised by forestry professionals who were experienced at conducting state forest Stand Level Inventory cruises. Additionally, the Private Forests monitoring team performed quality checks of every site each time a vegetation plot cruise was completed. We therefore believe these data are credible.




Figure 17. Predicted temperature increases from simulated FPA harvest of all sites. The leftmost boxplot, “All,” represents the mean response of all sites. The two dashed lines above 1 °C represent the 95% credibility interval for All, indicating the 95% probability that the mean response would be within that band. The other dashed line indicates 0.3 °C.

2.2.4 Percent Harvest

Prior to developing and testing prescriptions we wished to test the model’s performance against a suite of conditions. Since the model relies on changes in the variable *PctBasalAreaReduced* we performed an examination of anticipated temperature increases with incremental changes in this variable. The predicted temperature increase changed as each site’s basal area was reduced in 10% increments from 80% to 10% of pre-harvest values (Figure 19). Although the sites exhibit different temperature responses to basal area removal (Figure 19A) it appears that, when examining basal area removal within 100 ft horizontal distance of the stream, the 0.3 °C threshold is generally crossed at around a 15% removal of basal area (orange line in 19A, mean line in 19B). The mean lines indicate a 50% probability that the mean response will be below or above that line at the given level of basal area reduction. The credibility interval for the mean (Figure 19B) indicates that the probability of the mean response crossing 0.3 °C lies between a 12 and 18% basal area reduction.




Figure 18. Square feet of basal area from 170 ft x 500 ft plots adjacent to streams. The pre-harvest and as-harvested results were measured in the field. The FMP and FPA harvests were simulated harvests of the pre-harvest data. All private sites are to the right of the orange line; state forest sites are on the left. Vertical lines indicated the difference between as-harvested and the potential harvest level at a site.

2.2.5 Harvest by distance from stream

The percent basal area reduction scenario demonstrated that a small fraction of basal area removed from within 100 ft of the stream would cause temperature increases above 0.3 °C. We examined how reduction in basal area as a function of distance from stream would affect temperature change. Due to differences in distance measure between state forests and private timberlands, the harvest was conducted according to horizontal distance (Figure 20) and slope distance (Figure 21). Slope distance is either equal to or greater than horizontal distance. Figure 21 includes slope distances out to 120 ft as this captures all but 1.5% of trees within 100 ft horizontal distance of the stream. The other trees had slope distances >120 ft at horizontal distances of 100 ft.

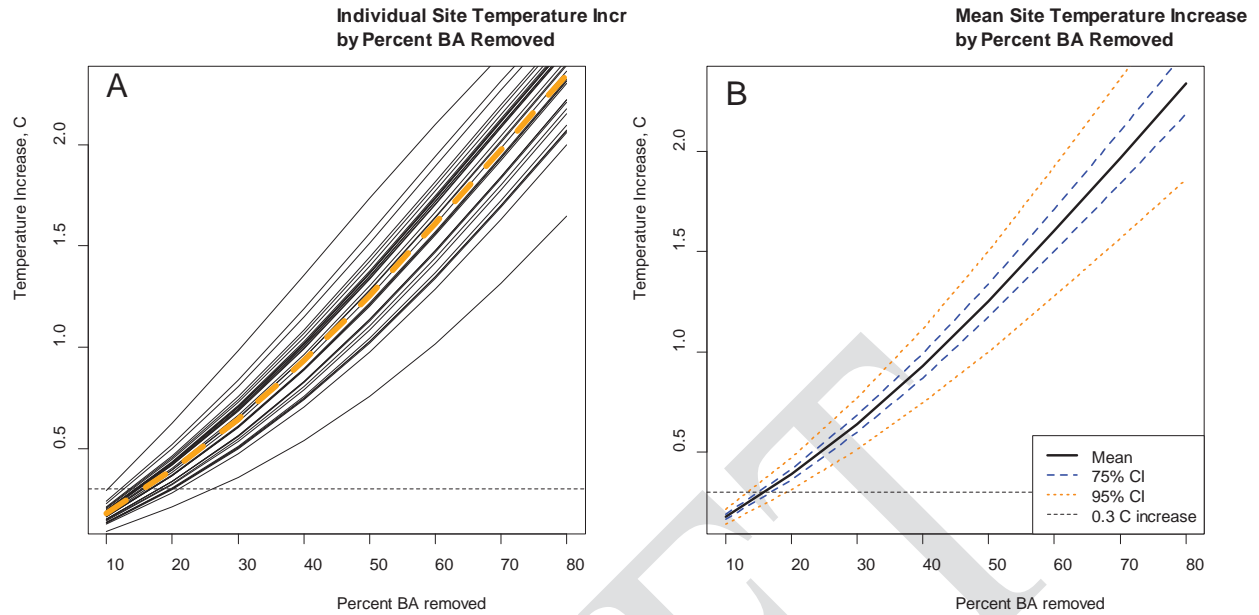


Figure 19. Temperature increases at a site level (A) and for an overall mean (B). The dashed horizontal line represents 0.3 °C. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the percent basal area removed from each site. The Y axis is the temperature increase (°C) due to the simulated harvests.

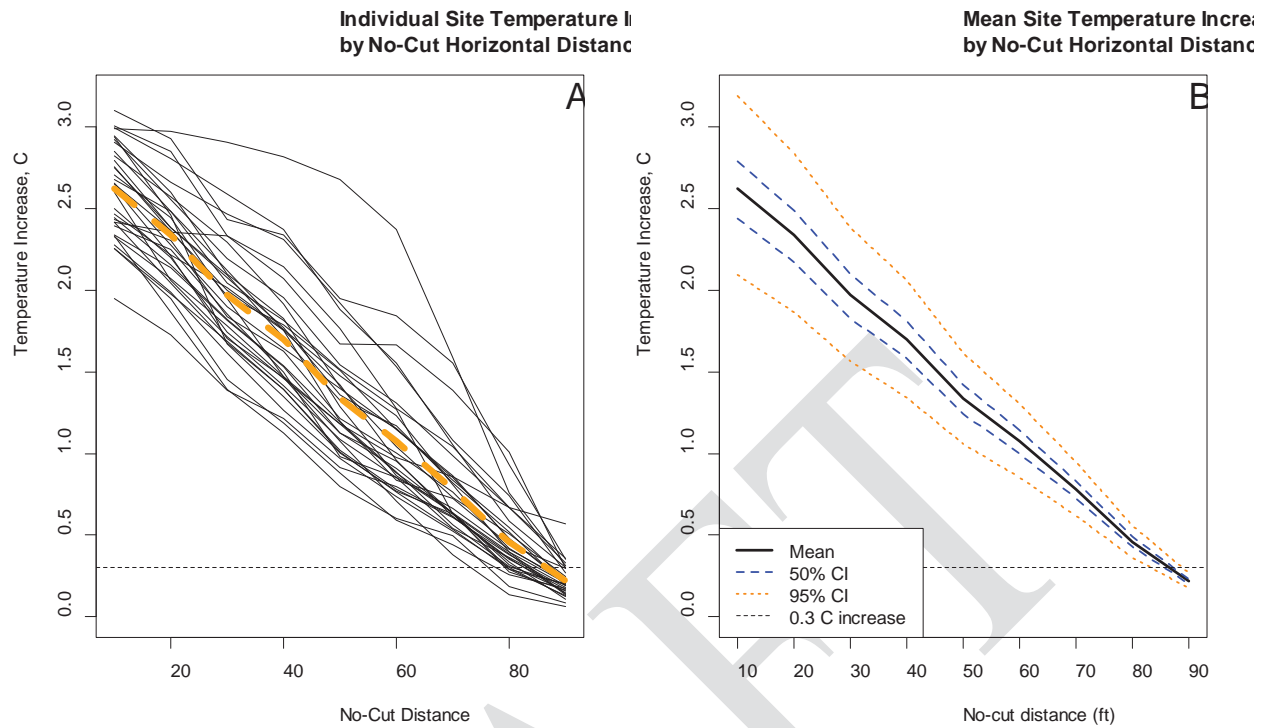


Figure 20. Predicted temperature increase at a site level (A) and overall mean (B) for harvest beyond specified horizontal distances. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the no-cut distance that was not harvested for both banks of every site. The Y axis is the predicted temperature increase (°C) due to the simulated harvests.

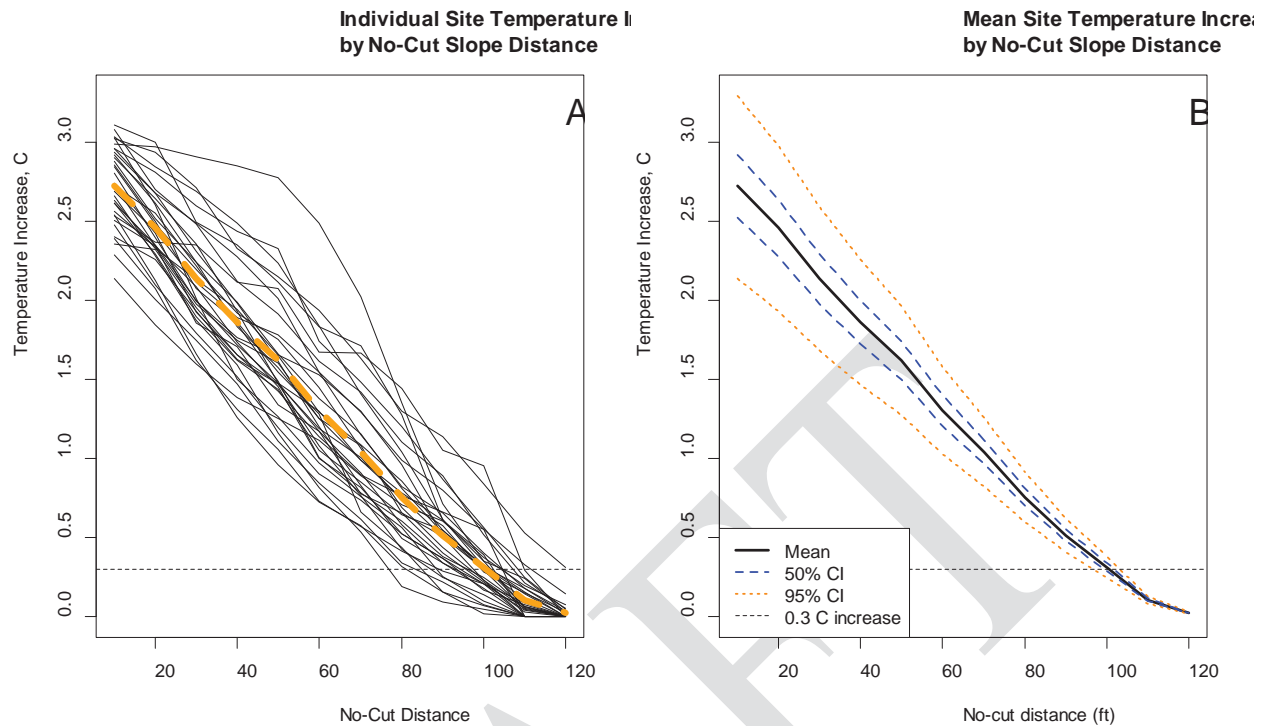


Figure 21. Predicted temperature increase at a site level (A) and overall mean (B) for harvest beyond specified slope distances. The orange line in A is equivalent to the line of the mean response in B. The blue and orange lines in B represent respectively the 75% and 95% Credibility Interval. The X axis for both graphs represents the no-cut distance that was not harvested for both banks of every site. The Y axis is the predicted temperature increase (°C) due to the simulated harvests.

Literature Cited

Ellison, A. M. 1996. An introduction to Bayesian inference for ecological research and environmental decision-making. *Ecological Applications* 6:1036-1046.